Analyzing the Effectiveness of Content Delivery Network Interconnection of 3G Cellular Traffic

Giyoung Nam
Department of Electrical Engineering, KAIST
giyoung@ndsl.kaist.edu

KyoungSoo Park
Department of Electrical Engineering, KAIST
kyoungsoo@ee.kaist.ac.kr

ABSTRACT

It has been a recent trend of Internet service providers (ISPs) to deploy content delivery networks (CDNs) extensively in their infrastructure in order to utilize their network resources and generate a new profit source. This 'telco CDN' has become an attractive solution since it enables the ISP to use their own network topology and real-time traffic information to address the bottlenecks, and seek for near-optimal path to convey the content to the users. However, since the location of telco CDN cache is limited to their ISP region, it became difficult to bring its performance benefit to the users outside the ISP region and it also led to suboptimal traffic reduction at ISP borders.

CDN interconnection (CDNi) is an emerging technology which has a potential to eliminate the redundant HTTP traffic received from external CDNs. A telco CDN can minimize the CDN traffic crossing the ISP border and at the same time deliver the content to their users quickly from its local cache by temporarily caching the content owned by a collaborating peer CDN. In this paper, we have studied the performance of CDNi when applied to the fast-growing cellular Internet traffic. We have simulated the CDNi protocol to gauge the bandwidth savings along with request redirection overheads using 7.7 billion HTTP logs (290 TBs by the byte volume) from one of the largest cellular ISPs in South Korea. We observed that 69% of total downlink traffic passes the Internet Exchange point (IXP), and according to our simulation results, intra-ISP CDN with CDNi can remove 16.2% to 29% of the IXP traffic. We also saw that the CDNi request redirection overhead could be significant to small objects, but it is still expected that if only large HTTP objects are redirected, a large bandwidth would be saved.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network monitoring

General Terms

Design, Measurement, Performance

Keywords

CDN, CDN interocnnection, IXP, Measurement

CFI'14, June 18–20, 2014, Tokyo, Japan

http://dx.doi.org/10.1145/2619287.2619297.

1. INTRODUCTION

Modern global Internet traffic increases rapidly with the growing popularity of high-resolution videos and smartphone applications. Cisco [1] predicted that the amount of global IP traffic increases by a factor of 3 in the next five years, and the aggregate traffic per year will reach a zettabyte by the end of 2017. Many Internet service providers (ISPs) are deploying content delivery networks(CDNs) to serve popular contents from their local cache to utilize their network resources with the fast traffic growth. This 'telco CDN' is attractive since the ISPs can exploit their own domain knowledge (such as detailed network topology and real-time traffic information) directly into CDN deployment and optimize the content delivery by dynamically adjusting the path to varying network loads. Many major ISPs such as AT&T, BT, Level 3, and Verizon have already been providing the CDN service [2] and even regional ISPs such as Korea Telecom and SK Telecom in South Korea are actively preparing the CDN service.

The goal of a telco CDN is twofold. One is to efficiently utilize the network infrastructure by reducing redundant data transfers spanning across their networks. By caching popular contents close to the downloaders, the ISP can not only reduce the network resource usage but also improve the response time to the users. Also, local CDN cache can potentially cut down the interdomain traffic crossing the Internet Exchange point (IXP) by absorbing the content delivery needs of the content providers that are outside of the ISP. This would increase the available bandwidth capacity of IXP and allow more efficient bandwidth usage at an ISP border. The second goal is to generate a new source of revenue in a cost effective manner. Since a telco CDN can provision the service without the bandwidth cost, it can significantly lower the cost of the CDN service.

One challenge of a regional telco CDN lies in the limited service coverage. Unlike existing CDN vendors that have global presence [3–6], regional telco CDNs may have to place their caching servers in their own networks. If a user from another ISP accesses the content hosted by a telco CDN, the content has to be delivered across the IXP with an extra delay compared with serving it to a local user. This would limit the performance benefit of the telco CDN only to its own subscribers, and decrease the opportunities to reduce the redundant traffic at IXPs.

To address this problem, many telco CDNs (as well as regional CDNs) with a limited service coverage are actively working on collaborative cache with a protocol called CDN interconnection (CDNi) [7–9]. CDNi allows temporarily sharing the popular content from participating peer CDNs located in different AS domains, and can potentially eliminate the CDN traffic crossing the IXP. It dynamically redirects the client request to a CDN close to the user such that the content is served from its local cache. In most cases, the serving CDN (e.g., the CDN that receives the redirected request) and the client are co-located in the same ISP, which effectively reduces the IXP traffic and better utilizes the internal network resources.

CDNi is still in its early stage, so there have been few studies on its effectiveness in the real-world traffic. In this paper, we analyze the benefit of CDNi with fast-growing cellular Internet traffic. We choose cellular Internet traffic since it is rapidly increasing and could benefit heavily from CDN and CDNi. We have logged 7.7 billion HTTP transactions at a 3G core network of one of the largest cellular ISPs in South Korea for seven days, and simulate the CDNi protocol to gauge possible IXP traffic reduction as well as the cost of CDNi request redirection. Our dataset covers 1.8/4.5 million unique client/server IP addresses respectively, and the total byte volume amounts to 290 TBs.

Our Web cache simulation shows that CDNi could be very effective in cellular Internet traffic. We find that 69% of the total traffic is coming from external ISPs and 16.2% to 29% of them can be reduced if we use collaborative CDN cache with CDNi. We also estimate the CDNi request redirection overhead, and find that serving the small objects with CDNi could produce noticeable delays. However, if we apply CDNi only to large objects, the overhead becomes reasonable with significant IXP bandwidth savings.

The rest of the paper is organized as follows. First, we provide the brief background of the CDNi protocol and explain the benefits of CDNi in detail in Section 2. Section 3 describes the cellular traffic dataset we use for our study, and Section 4 presents the Web cache simulation results from the perspective of ISPs and end users. Section 5 shows related works and we conclude in Section 6.

2. CONTENT DELIVERY NETWORK INTERCONNECTION

In this section, we provide brief background of the CDNi protocol and discuss the benefits from the perspectives of ISPs, content providers, and end users.

2.1 CDNi protocol

Figure 1 shows the basic request routing scheme of CDNi. CDNi request routing determines a CDN server from which a user downloads contents. CDNi defines two ways to determine the server. One is to use DNS redirection, which collaboratively resolves the IP address of the proper server with a DNS lookup. The other approach is to use HTTP redirection, which redirects the user's HTTP request to the appropriate peer CDN server close to the user. We focus on HTTP redirection here since it allows more fine-grained request redirection policies.

HTTP redirection of CDNi is very simple. When a user in one ISP sends a request to the original CDN (called *upstream CDN*) for content access, the upstream CDN decides if the request should be redirected to a closer CDN server from the location of the requestor. If it finds a better CDN (called *downstream CDN*) than itself (e.g., a peer CDN located in the same AS that hosts the user's IP address), it responds with an HTTP 302 redirection response with the location of the downstream CDN servers. When the user sends the request to the new location, the downstream CDN can further redirect it to another peer CDN or decide to deliver the content directly to the user. The process of request redirection can always involve the user (called *iterative redirection*) or can transparently resolve the final CDN server location among the collaborative CDNs and send only the destination server to the user (called *recursive redirection*). If the content is a cache miss at the downstream CDN server, the server fetches the content from the upstream CDN, caches it, and delivers it to the user. In CDNi involving telco CDNs, an upstream CDN in one ISP typically redirects the requests to a local CDN to the users, which helps reduce the IXP traffic as well as improve the user response time.

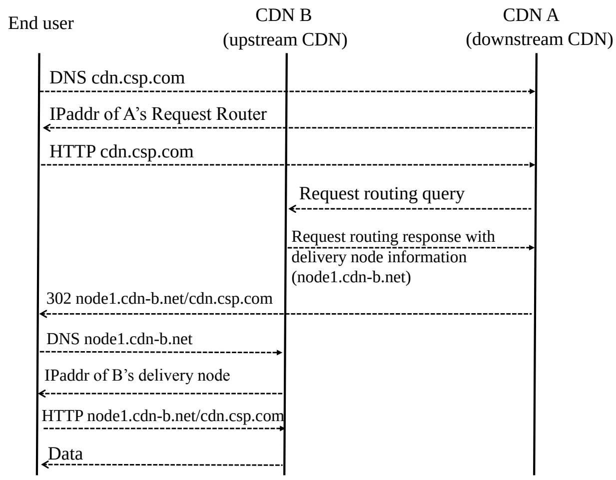

Figure 2: **An example of CDNi redirection protocol (HTTP).**

DNS redirection is much simpler than HTTP redirection. Unlike the HTTP redirection protocol sends 302 response from CDN A, CDN A sends a CNAME response while a DNS query comes from end user. DNS redirection reduces 1 RTT delay. As shown in Figure 2, the end user does not need to send an HTTP query to CDN A. However, DNS-based redirection can see the DNS resolver's IP address instead of the end user's IP address during redirection. That is, DNS redirection enforces coarse-grained request redirection policies.

In addition to request routing, the CDNi specification defines control, metadata, and logging interfaces to exchange the service information among collaborating CDNs [7]. The control interface first bootstraps other interfaces, and communicates with peer CDNs for content pre-positioning, revalidation, and purgation of metadata and content fetched from peer CDNs. The metadata interface exchanges content metadata such as authentication information and content type, and the logging interface manages content and flow logs, and designates a CDN for accurate log processing which includes log analysis and accounting for billing. Each CDN has a CDNi gateway that implements control, metadata, and logging interfaces and exchanges the CDNi information with peer CDNs.

2.2 Benefits of CDNi

By deploying the CDNi, one can expect to overcome some of the limitations of regional telco CDNs. We categorize the CDNi benefits from the viewpoints of ISPs, content providers (CP), and end users, respectively.

Benefits to ISPs Essentially, CDNi allows sharing the resources of collaborating peer CDNs to extend the service coverage of a regional CDN. By temporarily caching the contents that originally belong to external CDNs, a regional telco CDN can serve the contents from its local cache to the users in the same ISP region. This should help reduce the traffic by the external CDN that crosses the IXP traffic. Also, from the perspective of the upstream CDN, CDNi extends the service coverage to the outside of its own region. For example, if a user in the U.S. accesses a content hosted by a regional telco CDN in South Korea, the Korean CDN can redirect the user request to a peer CDN in the U.S. and have it deliver the

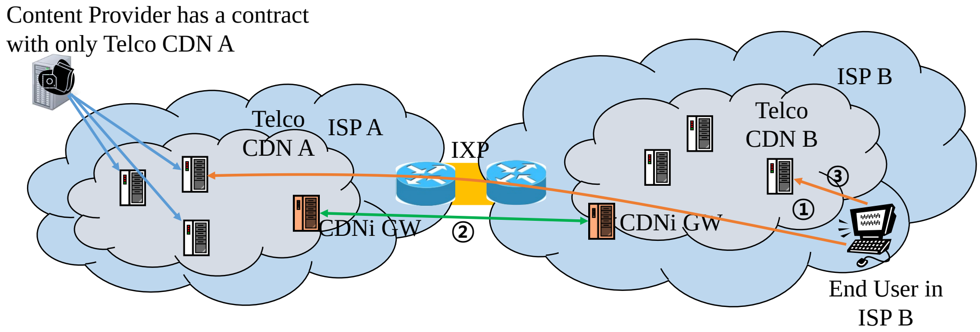

Figure 1: A CDNi request routing scenario. We assume that the content is already cached at a telco CDN of ISP B. (1) an end user send a request to the original CDN in ISP A, and (2) CDNi gateways exchange user request information to select a proper CDN. After the end user get a redirection message from CDN A, (3) the user downloads the content from the selected CDN which is close to the user.

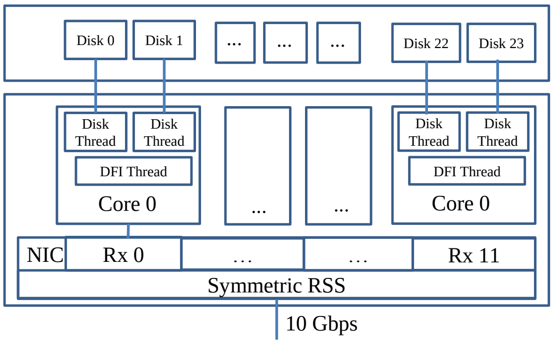

Figure 3: Overall architecture of Monbot

content locally to the user. This should allow efficient multiplexing of many collaborating CDN resources, which is also found in other types of multiplexing like cellular data roaming.

Benefits to CPs From the perspective of content providers, CDNi could potentially lower the cost of CDN service. Traditionally, if a CP wants to distribute their contents to many geographic regions, it has to make a contract with a CDN vendor that has global presence, which can be burdensome to small CPs. CDNi allows a small CP to make a contract with a regional CDN that participates in CDNi, and benefits from the global CDN service at a low cost.

Benefits to end users End users could benefit from fast response from CDNi for contents that are originally served from the outside of the user's ISP. CDNi makes the location of the contents local to the users, and the regional telco CDN can even dynamically optimize the delivery path considering the user's location and network load situation.

Despite these benefits, there exist few studies that analyze the effectiveness of CDNi with real-world traffic. We set out to look into it with fast-growing cellular Internet traffic in this paper, and see how much traffic can be saved by a local CDN and how much extra bandwidth savings CDNi can bring. We also analyze the overhead of CDNi by estimating the request redirection delays and suggest peering strategies for better network resource utilization.

3. MEASUREMENT ENVIRONMENT AND DATASET

In this paper, we analyze the HTTP traffic at one of the largest cellular ISP in South Korea, obtained from our prior work [10]. We log all HTTP request and response pairs that pass one of three 10 Gbps core network links just below a 3G GPRS support node (GGSN) in Seoul for one week (from 11am on July 7th to 2pm on 13th in 2012). We use the high performance deep flow inspection system (DFI) called 'Monbot'. This system produces three types of log data in real time; 1) TCP flow statistics, 2) HTTP request/response statistics, and 3) SHA-1 hashes of content chunks. This system runs on commodity hardware and captures packets in the high-speed 10Gbps network without packet drop. Monbot employs RSS to assign the packets which belong to the same TCP connection to the same CPU core. Each core processes these packets and writes the log into two disks to avoid disk I/O bandwidth bottleneck.

The cellular ISP has 12.5 million subscribers, and the measurement point covers the half of the regions served by the ISP in South Korea. During the measurement period, we have logged 7.7 billion HTTP requests which amount to 290 TBs in byte volume, taking up 75% of total downlink traffic. Specifically, we log the HTTP request and response headers as well as TCP connection time, connection setup delays, and source and destination IP addresses and port numbers. The numbers of unique client and server IPs that we saw are 1.8 and 4.5 million, respectively. The reason why the number of unique client IP addresses is smaller than that of servers is because the cellular ISP dynamically assigns a private IP address to each user device, and different devices may reuse the same IP address.

From this dataset, we determine the incoming traffic that crosses the IXP based on the server IP address. We use a GeoIP database from Maxmind [11] to map a server-side IP address to its hosting ISP and country. For Web cache simulation, we follow the HTTP 1.1 standard [12], and assume that a CDN caching server uses a log-based disk cache as in [13]. We also calculate flow size distribution, average bandwidth, and check the connection setup delays and TCP flow duration to estimate the overheads of CDNi request routing and the impact on the performance for end users.

4. EFFECTIVENESS OF CDNI IN CELLULAR INTERNET TRAFFIC

In this section, we discuss the benefits and overheads of CDNi with the real-world 3G cellular traffic.

4.1 Bandwidth Savings by CDNi

We first measure how much of the downlink traffic is coming from the outside of the cellular ISP, and simulate Web caching to

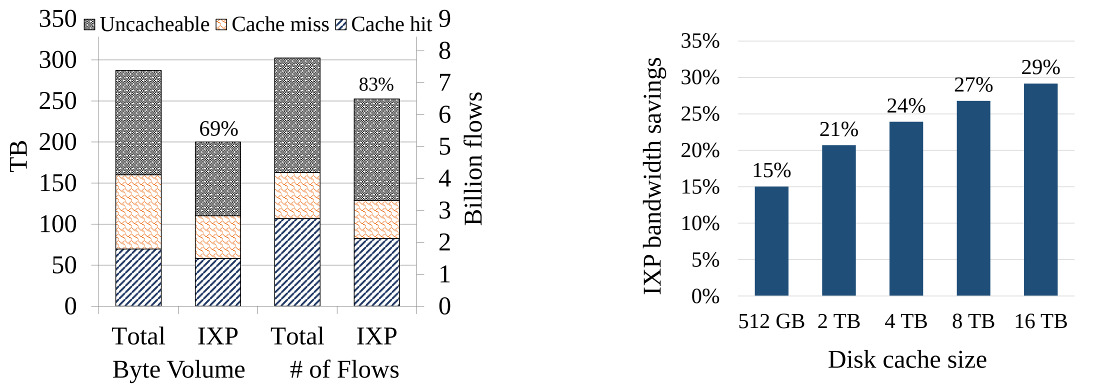

(a) Breakdown of Web cache hit/miss objects of total and IXP traffic with 16 TB disk cache

(b) IXP bandwidth savings over various Web cache sizes

Figure 4: Web cache simlation and IXP bandwidth savings for various cache size

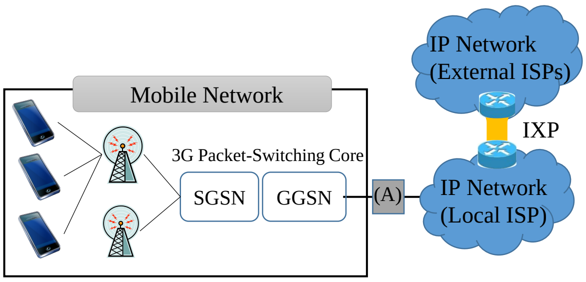

Figure 5: Overall 3G mobile network structure. (A) is the location of Web cache middlebox in simulation.

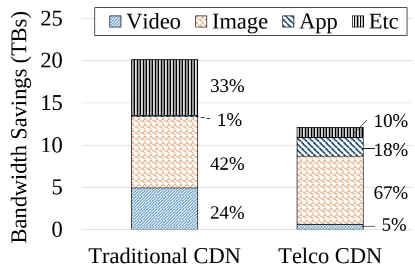

Figure 6: Breakdown of IXP bandwidth savings by popular content types with external CDNs

see potential IXP bandwidth savings with CDNi. Since we cannot predict how many external CDNs will participate in CDNi with this ISP, we explore two extreme points here. (a) How much IXP bandwidth savings do we expect if all servers outside the ISP are hosted by CDNs that participate in CDNi, and (b) what will be the IXP bandwidth savings if all traffic currently served by external CDNs is handled by CDNi. (a) shows the maximum bandwidth savings this ISP can achieve with the current traffic mix and (b) presents more practical estimates with the current set of external CDNs if this ISP runs the CDNi protocol with them.

To probe (a), we run the Web cache simulation with our 3G mobile HTTP request log. Figure 5 shows the architecture of the 3G mobile network. In the Web cache simulation, we assume that a transparent caching middlebox is placed at (A) in Figure 5, so that all HTTP requests that we have collected form the logs pass through the Web caching middlebox. Our web cache simulator assumes log-based disk cache, and it determines cacheability of a Web object based on 'max-age' field in HTTP response and caches it. If response is special type codes such as '304 Not Modified', or there is a prohibit message in the response header such as 'no-store' or 'no-cache', Web caching middlebox sets the object to 'uncacheable'.

Figure 4(a) shows the breakdown of Web objects when we simulate Web caching with 16 TB disk cache. We first note that the majority of the incoming traffic passes the IXP of this ISP. We see that 83% of the total downlink flows (or 69% by the byte volume) comes from external ISPs. This implies that the object size crossing the IXP is smaller (30.82 KB on average) than those served locally in the ISP (60.77 KB on average). We will discuss the reason for this later in this section. We find that as much as 58 TB of the IXP traffic (out of 200 TB that crosses the IXP) can be saved and served locally by CDNi. This shows that 29% of the IXP traffic can be removed if we assume that all external servers participate in CDNi. Figure 4(b) shows the effectiveness of Web caching for the IXP traffic over various cache sizes. Not surprisingly, the amount of bandwidth savings increases as the Web cache size grows, but the benefit quickly diminishes over large sizes. If we use an infinite cache, 31% of the IXP traffic can be saved by CDNi.

Since the assumption that all external servers are host by CDNs is rather optimistic, we go on to measure more realistic bandwidth savings by applying CDNi only to the interdomain traffic hosted by current CDNs. We extract the flows of external CDNs by matching the server domain names against the CDN names listed in [2]. We find that 32% of the IXP traffic is served by traditional non-telco CDNs, and 17% of the IXP traffic comes from external telco CDNs. Overall, 16.2% of the traffic by external CDNs (e.g., 20 TB of traditional CDN traffic, and 12 TB of telco CDN traffic) can be saved if they participate in CDNi with this ISP. Figure 6 shows the IXP bandwidth savings by popular content types, and Table 1 shows the breakdown of the flows and byte volumes divided into the locations of the servers.

Finally, we look into the internal ISP traffic served by CDN servers located inside the ISP. Surprisingly, Google Cache alone handles 48.7% of the total intra-ISP traffic. We find that the majority of the traffic is due to YouTube and Google Play that provide

Servers	Total		Video		Image		App	
	# Flows	Volume (TB)	# Flow Ratio(%)	Volume Ratio(%)	# Flow Ratio(%)	Volume Ratio(%)	# Flow Ratio(%)	Volume Ratio(%)
Total	7,790,506,152	287.63	0.59	32.31	48.82	28.87	18.26	22.43
Total IXP	6,489,701,334	200.02	0.50	32.77	47.55	33.06	18.43	18.83
External CDN (IXP)	1,534,229,820	65.77	0.93	39.86	59.16	28.66	8.82	10.54
Telco CDN (IXP)	1,068,149,157	34.87	0.06	13.22	7.14	46.39	1.35	19.53
Total Local ISP	1,300,804,818	87.61	1.02	31.28	52.90	19.27	17.41	37.60
Google Cache	12,895,334	42.43	62.43	49.70	< 0.01	< 0.01	28.71	50.28
Local CDN	464,029,373	12.35	0.41	18.04	81.59	57.32	5.77	7.79

Table 1: Breakdown of the traffic by popular content types and locations of servers

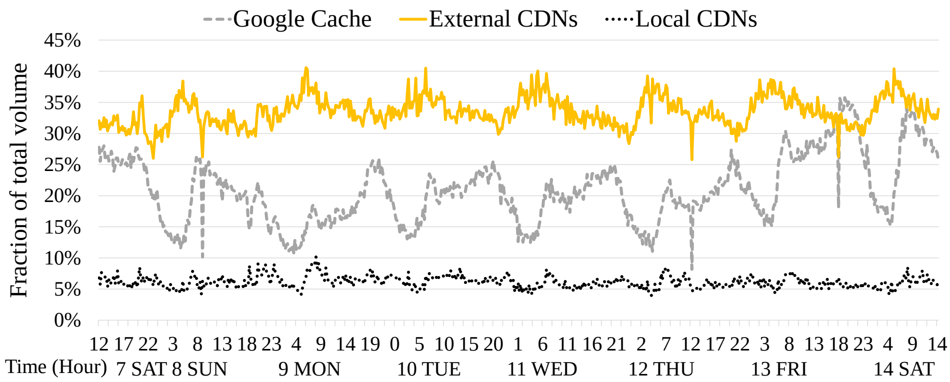

Figure 7: Fraction of traffic volume by Google Cache, external and local CDNs for seven days

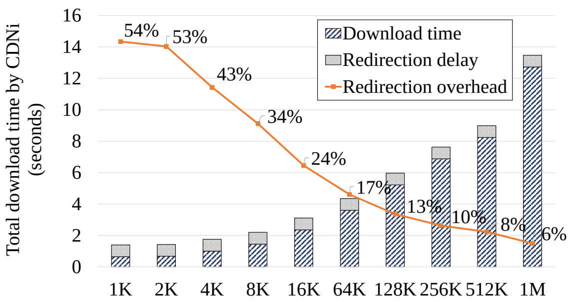

Figure 8: CDNi request redirection overhead in HTTP object download times

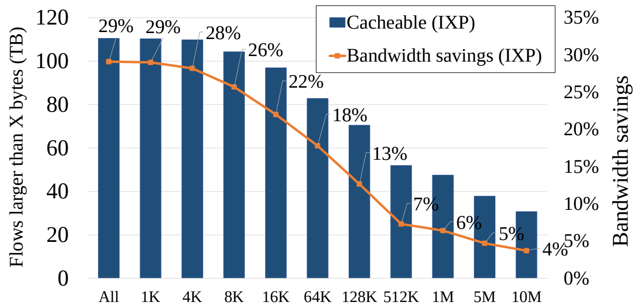

Figure 9: Bandwidth savings of caching only the objects larger than X bytes

video streaming and smartphone application binaries. This explains why the average object size of intra-ISP flows is larger than that of external ISP flows. Other local CDNs handle 14.2% of the intra-ISP traffic, so 62.9% of the total intra-ISP traffic is served by local caching servers. Figure 7 compares the byte volumes consumed by Google Cache, external, and local CDNs over the entire measurement period. While Google Cache takes up a significant portion of the entire traffic, its byte volume ratio fluctuates over time. In contrast, the traffic by the CDNs takes up more or less constant portion of the entire traffic. This implies that bandwidth multiplexing of traffic by multiple CPs works particularly well with CDNs while Google Cache may have to overprovision their resources for the peak time.

4.2 Request Routing Overhead of CDNi

CDNi works by redirecting the request to a downstream CDN, which adds one HTTP redirection delay to the user. For small HTTP objects, this extra delay could be significant in the object download time, and could reduce the quality of experience (QoE) of the users. We estimate the request redirection overheads by calculating the average response time of HTTP 302 redirection messages of external CDN servers. We see 0.74 second as the average redirection overhead. We plug this number into the object download time to estimate the CDNi overhead. Figure 8 shows the percentage of the CDNi request routing delay in the download time of HTTP objects served by intra-ISP CDNs. We find that the redirection overhead is severe with small objects(e.g., 1 KB objects), but the overhead decreases as the object size increases. At 8 KB objects, the CDNi request redirection overhead is reduced to 34%, which may be reasonable for user QoE.

We explore how much bandwidth savings an ISP expects if CDNi limits the request redirection to only large objects. Figure 9 shows the IXP bandwidth savings when we redirect only the objects that are larger than X bytes. Interestingly, the IXP bandwidth savings by the objects smaller than 4 KB is negligible (less than 1%). This is because the aggregate byte volume that these objects take up is small compared to that of larger objects though the number of such objects takes up the majority of HTTP flows (~60%). Even if upstream CDNs redirect only the objects larger than 8 KB, we can expect a significant IXP bandwidth savings (∼25.7% out of 29% maximum bandwidth savings).

5. RELATED WORKS

There have been many works that attempt to reduce the IXP traffic by peer-to-peer (P2P) systems [14–16]. P4P exposes explicit cooperation between P2P systems and ISPs to fetch the content from an intra-ISP peer with a minimal hop count [14], and [15] utilizes a global CDN's redirection information for peer selection. [16] inserts ISP-owned peers to reduce the BitTorrent traffic crossing the IXP. While these techniques are effective in reducing the cross-domain P2P traffic, recent trend shows that the P2P traffic is decreasing with the popularity of streaming videos by CDNs [17]. Also, the P2P traffic in the cellular Internet traffic is small [10], so the effectiveness of previous techniques could be limited in the cellular traffic.

The HTTP traffic takes up the majority of bandwidth consumption (∼75%) in the current 3G and LTE traffic [10,18]. While Web caching is effective in the cellular traffic, transparent network redundancy elimination (NRE) on the TCP flow level could produce more bandwidth savings. [10] shows that one can expect a bandwidth savings of up to 59% of the entire 3G downlink traffic even if the ISP uses a simple NRE scheme with 4KB fixed-size chunking. In contrast, we gauge the effectiveness of standard Web caching from the perspective of CDNi. CDNi is much simpler than NRE and can be more easily deployed with existing HTTP CDNs.

High speed packet capturing is the first step of our work, and there are many works that technique to capture high speed network packets in commodity NICs [19–21]. PF_RING [20] allocates a large packet ring buffer in driver (kernel space) and batch the packet from NIC to ring buffer. PF_RING also allocates intermediate data buffer called 'pf_ring' within the kernel and uses it as data capture buffer for each receive ring. Netmap [21] is similar with PF_RING, but application can directly access the ring buffers in kernel space using shadow copy, and it reduces extra memory copy. Monbot use the PacketShader I/O engine [19] to process all packet in the same TCP connection in the same core by RSS, so each application thread can read packets directly from multiple RX queues in NIC.

6. CONCLUSION

CDNi is an emerging technology that extends the CDN service coverage of regional telco CDNs. In this work, we have presented practical benefits and overheads of CDNi with the cellular Internet traffic. Our study shows that one can expect as much as 16.2% IXP bandwidth savings if all current CDNs participate in the collaborative caching. We also find that the request routing delay overhead of CDNi could be significant to small objects, but if upstream CDNs redirect the objects larger than 8 KB only, CDNi expects a large IXP bandwidth savings since small objects do not take up much of bandwidth consumption. For the current traffic mix, images and videos would benefit from CDNi the most, and as the traffic by large objects increases, the benefit from CDNi will grow in the future.

7. ACKNOWLEDGEMENTS

We appreciate many valuable comments on the earlier version of this paper from anonymous reviewers of CFI 2014. This research was funded by the MSIP (Ministry of Science, ICT & Future Planning), Korea in the ICT R&D Program 2014 (Project no. 2014-044-008-001).

8. REFERENCES

[1] CISCO. Cisco Visual Networking Index: Forecast and Methodology, 2012 to 2017. Technical report, 2013.

[2] Telco CDN lists. `http://en.wikipedia.org/wiki/Content_delivery_network#Telco_CDNs`.

[3] Akamai Inc. `http://www.akamai.com/`.

[4] Limelight Networks Inc. `http://www.limelight.com/`.

[5] Amazon CloudFront. `http://aws.amazon.com/cloudfront/`.

[6] CDNetworks Inc. `http://cdnetworks.com/`.

[7] Content Delivery Networks Interconnection. `http://tools.ietf.org/wg/cdni/`.

[8] B. Niven-Jenkins, F. Le Faucheur, and N. Bitar. Content Distribution Network Interconnection Problem Statement. `http://tools.ietf.org/html/rfc6707`.

[9] G. Bertrand, E. Stephan, and T. Burbridge. Use Cases for Content Delivery Network Interconnection. `http://tools.ietf.org/html/rfc6707`.

[10] S. Woo, E. Jeong, S. Park, J. Lee, S. Ihm, and K. Park. Comparison of Caching Strategies in Modern Cellular Backhaul Networks. In *Proceedings of the 11th International Conference on Mobile Systems, Applications, and Services (ACM MobiSys'13)*, 2013.

[11] Maxmind GeoIP database. `http://www.maxmind.com/`.

[12] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-lee. Hypertext Transfer Protocol - HTTP/1.1. `http://tools.ietf.org/html/rfc6707`.

[13] A. Badam, K. Park, V.S. Pai, and L.Peterson. HashCache: Cache Storage for the Next Billion. In *In Proceedings of the USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2009.

[14] H. Xie, Y.Richard Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz. P4P: Provider Portal for Applications. In *Proceedings of the ACM SIGCOMM conference on Data communication*, 2008.

[15] David R. Choffnes and Fabian E. Bustamante. Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems. In *Proceedings of the ACM SIGCOMM conference on Data communication*, 2008.

[16] I. Papafili, S. Soursos, and G. D. Stamoulis. Improvement of BitTorrent Performane and Inter-domain Traffic by Inserting ISP-Owned Peers. In *Network Economics for Next Generation Networks*, 2009.

[17] D. Weller and B. Woodcock. Internet Traffic Exchange: Market Developments and Policy Challenges. *OECD Digital Economy Papers*, No. 207, 2013.

[18] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck. An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*, 2013.

[19] S. Han, K. Jang, K. Park, and S. Moon. PacketShader: a GPU-Accelerated Software Router. In *Proceedings of the ACM SIGCOMM conference*, 2010.

[20] PF_RING. `http://www.ntop.org/products/pf_ring/`.

[21] L. Rizzo. Netmap: A Novel Framework for Fast Packet I/O. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC'12)*, 2012.