Towards Understanding Developing World Traffic

Sunghwan Ihm
Computer Science
Princeton University
sihm@cs.princeton.edu

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

Vivek S. Pai
Computer Science
Princeton University
vivek@cs.princeton.edu

ABSTRACT

While many projects aim to provide network access to the developing world or improve existing network access, relatively little data exists regarding the behavior of traffic in these environments, especially in regards to the characteristics of traffic in the developing world. In this paper, we provide a first glimpse into the traffic gathered by a worldwide proxy network, and try to observe differences in first-world and developing-world traffic characteristics. What sets this work apart from similar research is the scope and level of detail – we capture more than 3TB of content representing one week's browsing by 348K users across 190 countries. Capturing the content, rather than just access logs, also allows us to perform similarity analysis at the content level.

Categories and Subject Descriptors

C.2.m [**Computer-Communication Networks**]: Miscellaneous

General Terms

Measurement, Performance

Keywords

Web Caching, Traffic Redundancy, Developing Region

1. INTRODUCTION

Internet access in the developing world is often a scarce commodity, and when it does exist, its costs in absolute terms can far exceed the cost of similar capacity in the first world. As such, techniques such as Web caching [6, 13] are commonly used to reduce the traffic demands of these environments, but it may be difficult to gauge the efficacy of these approaches as the Web evolves and traffic characteristics change. In particular, Web proxies cache entire objects (HTML text, images, etc.) that are static and do not change on a per-user basis. As such, Web proxies typically provide little benefit for customized content or for content that changes often relative to its access frequency. Some measurements of Web usage in the developing world have been performed by Du et al. [8] for Cambodia and Ghana using Web proxy access logs, which provide an object-level view of Web traffic.

Our interest in developing-world usage is to understand the impact of the uncached content, since it is this content that must be fetched over the wide-area network (WAN) that connects the Web proxy to the rest of the Internet. To analyze the behavior of this content, we need access to the actual data being transferred, rather than just the access logs from a Web proxy. With a better understanding of the content, we can explore techniques to further reduce the wide-area traffic, using techniques to exploit duplicate content within objects [15, 16].

While some of these techniques have been shown to be effective in first-world environments, we lack detailed understanding of their effectiveness in the context of the developing world. Some questions we may want to answer are: How well do these techniques work? Are there regional variations in traffic that affect their behavior? Can chunk-level caching reduce cache miss traffic?

Answering these questions would allow us avenues for new opportunities to complement some of the research to improve connectivity in these environments [7, 18, 19, 23], and to better make use of frugal system that can operate using slow disks and limited memory [3, 4, 9]. We view content-focused analysis as a way to better understand the traffic optimization opportunities in these systems.

In this paper, we attempt to answer these questions by analyzing real Web traffic. In particular, we collect one week's worth of wide-area traffic in the CoDeeN [25] system, a heavily-used global Web cache.

We find several interesting results: (1) While developing-world users have poorer connectivity, this is not obviously reflected in the distribution of object sizes on the pages they visit, suggesting content transformation remains an open avenue for enhancing user experience, (2) canceled transactions, where a transfer is aborted mid-stream, account for a substantial portion of all traffic, and the rate of cancellation is actually higher in the first world than in the developing world, suggesting that generally faster response times are leading to user impatience, (3) developing-world users seek audio and video as a higher fraction of traffic than their first-world counterparts, and (4) the working set of content in the developing world is much larger than the size of typ-

NSDR'10, June 15, 2010, San Francisco, California, USA.

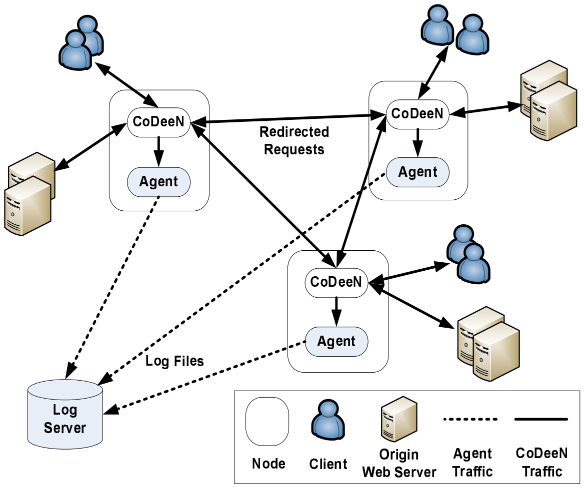

Figure 1: Data Collection

Capture Duration	3/25/2010 – 3/31/2010 (7 days)
Total # Requests	157,797,656
Total Bytes	3,132,405,067,559
Total # Clients	348,710
Total # Countries	190
/8 networks	157 (61.33%)
/16 networks	15791 (24.09%)
/24 networks	210947 (1.26%)

Table 1: Data Set Summary

ically deployed Web caches, and that much larger storage would provide a usable benefit.

The rest of this paper is organized as follows: we describe the details of data set we use in Section 2, and report our preliminary results in Section 3. Finally, we discuss related work in Section 4, and conclude in Section 5.

2. DATA SET

The traffic we use in our analysis comes from the CoDeeN content distribution network (CDN) [25], a semi-open globally-distributed proxy which has been running since 2003, and now encompasses more than 500 PlanetLab [20] nodes. This system serves over 30 million user-facing requests per day, routing traffic between CoDeeN nodes as appropriate. Some fraction of these requests are cache misses that must be retrieved from the origin Web servers that generate the content. We run an agent on each CoDeeN node that copies the traffic between CoDeeN and the origin servers, and sends it to a central collection server for processing. Figure 1 depicts our data collection methodology.

All of this content represents cache misses that could not be satisfied by either the client's browser cache, local proxy cache (if present) or by the CoDeeN proxy cache. In addition to logging the URLs of all requested content like a proxy cache would, we log all of the content itself. Each record in the log file contains the following information: client IP address, full request including all headers, downloaded object size, full response including its header and body, timestamp at the end of the transaction, and download time. Any difference between the downloaded size and the object's size is due to the client terminating the transfer, either explicitly or by moving to a different page.

We collect one week's traffic from March 2010, during which time we observe 157 million cache miss requests from 348K unique client IPs. The total volume of the traffic is about 3 TB, and the client population covers 190 countries, as determined using the MaxMind database [14]. The client IPs cover more than 60% of /8 networks, and 24% of /16 networks. Table 1 presents a summary of this data. To the best of our knowledge, this collection covers a client population several times larger than previous work [26], and is the only one that has captured the actual content (as opposed to access logs) at this level of detail.

	OECD	DevReg
Countries	27	163
Requests	25.68%	74.32%
Bytes	25.81%	74.19%
Clients	25.53%	74.47%

Table 2: OECD and DevReg Data Summary

3. PRELIMINARY RESULTS

In this section, we present initial results of our analysis, including country-level information, object analysis, site popularity, cacheability, and cancellation behavior.

To compare differences between the first world and the developing world, we divide the 190 countries in our logs into two groups. The first world ("OECD") consists of 27 high-income economies from OECD [17] member countries, excluding three upper middle-income economies, Mexico, Poland, and Turkey. We regard the remaining 163 countries as the developing world ("DevReg"). As shown in Table 2, OECD represents about 25% of the data set in most categories.

Table 3 shows the top countries by the number of requests, the total bytes, and the number of clients (unique IPs). In general, Poland, China, and Saudi Arabia are the heaviest users, appearing near the top in all of the categories. Middle-Eastern countries such as Saudi Arabia, Kuwait, and the United Arab Emirates are also heavily represented, accounting for 30% of the total byte traffic.

3.1 Object Size and Content Type

We first examine object-level traffic characteristics in these two data sets. The size distribution is shown in Figure 2 (a), and excludes outliers such as Poland and the Middle East. Per-country breakdowns by object type are also shown in Figure 2 for (b) text, (c) images, (d) video, (e) audio, and (f) application such as Flash. Not too surprisingly, we see that the median object size is small in the developing world, but even the OECD countries see most objects that are smaller than 10KB, consistent with previous traffic models [5, 12]. The difference between object sizes may reflect the importance of regional sites that are more tailored toward the bandwidths seen in the developing regions.

The per-object breakdowns provide some insight into this divergence. The percentage of responses for text objects (text and HTML) range from 30-90% for OECD countries and 10-100% for developing countries. The countries with low text fractions are largely browsing images, often blocked by local censorship. The high-text countries are likely due to a combination of poor bandwidth along with highly-cached images, which are not re-fetched. The image breakdown is largely the complement of the text breakdown – countries with a high text fraction have a low image fraction, and

Rank	Requests %	Bytes %	Clients %
1	Poland (23.67)	Saudi Arabia (19.35)	China (19.57)
2	China (15.24)	China (15.53)	Poland (13.73)
3	Saudi Arabia (11.51)	Poland (11.59)	Germany (8.70)
4	Germany (8.84)	United States (8.87)	Russian Federation (7.43)
5	United States (6.11)	United Arab Emirates (6.63)	Saudi Arabia (4.06)
6	Russian Federation (4.15)	Germany (6.52)	United States (3.84)
7	United Arab Emirates (4.02)	Kuwait (3.84)	Ukraine (3.05)
8	Kuwait (2.48)	Russian Federation (2.86)	Brazil (2.20)
9	Spain (2.13)	Spain (2.50)	Italy (2.00)
10	Ukraine (1.43)	Turkey (1.52)	France (1.90)

Table 3: Top Countries By Requests, Bytes, and Clients

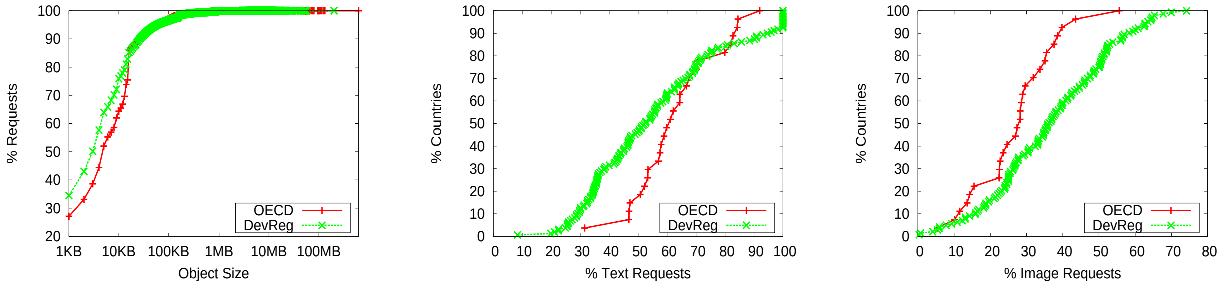

(a) Object Size CDF (1KB Bin) (b) % Text Requests CDF (c) % Image Requests CDF

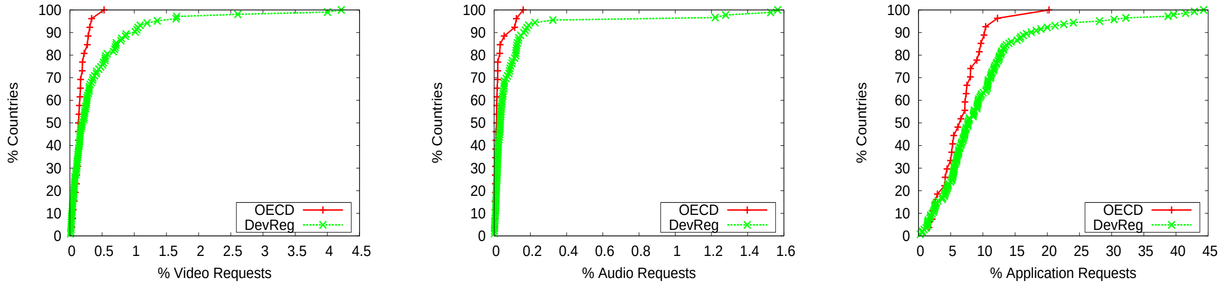

(d) % Video Requests CDF (e) % Audio Requests CDF (f) % Application Requests CDF

Figure 2: Object Size and Content Type

Rank	Requests %	Bytes %
1	Game (11.25)	Portal (5.43)
2	Game (11.00)	Video (5.43)
3	Portal (5.84)	Game (5.23)
4	Game (4.59)	Game (4.26)
5	Portal (2.56)	Game (2.69)
6	Portal (1.97)	Portal (2.59)
7	Ads (1.14)	Adult (2.58)
8	Chat (1.08)	Adult (1.88)
9	Analytics (0.82)	Adult (1.67)
10	Adult (0.77)	Portal (1.05)

Table 4: Top Sites By Requests, and Bytes

countries with very few text objects often download many thumbnailed images per HTML page.

The main surprises are seen in the audio and video categories, where the developing regions almost uniformly download more rich media as a fraction of their traffic when compared to the OECD countries. Were this purely a function of the same browsing that we see with images, we would expect a broader range of behaviors in the DevReg group, as seen on the graph for text objects. While the higher-bandwidth OECD countries no doubt consume more audio/video content, this data suggests that the DevReg group has a strong interest in rich media, despite their lower bandwidths. A significant fraction of this content appears to be music videos and songs in the MP3 formats.

The application objects, which include all Flash content, are also higher as a fraction of traffic in the DevReg group, and the absolute percentage of traffic is quite high, with medians near 7% of all requests, something that traffic modelers will have to consider for future work.

3.2 Site Popularity and Working Set Size

Using the combined traffic, we analyze the most popular sites and categorize them by content type in Table 4. While portal sites and game sites are popular, the presence of adult sites in the top 10 is somewhat surprising because previous work [8] measured adult traffic at roughly 1% both in terms of requests and bytes. We believe this may stem from usage, since they used logs from kiosks and Internet cafes, adding a public element to the browsing. Our data set presumably includes more private browsing at home, where such content may be more freely viewable.

We also measure the diversity of Web sites each country visits in Figure 4. We can see that the number of unique Web sites in the developing world is smaller than in the first world, typically numbering in the hundreds to low thousands. The OECD countries, in comparison, typically view

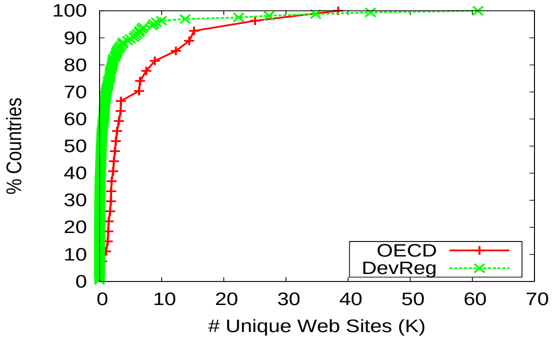

Figure 3: # Unique Web Sites CDF

a few thousand to tens of thousands of sites in this study. This may be a function of bandwidth, cost, or latency – randomly browsing sites is easier when browsing is faster and no extra expense is incurred.

3.3 Redundancy and Caching

Given that our accesses are all cache misses, we focus on using content-based caching [9,21,24] instead of object-based caching. This analysis attempts to find redundancy in content that is marked uncacheable at the HTTP level. Using 64 KB chunks searches only for coarse-grained redundancy, such as entire objects or parts within a large piece of content, like audio or video. This chunking finds about 30% redundancy consistently on a daily basis, starting from an empty cache. Using 1 KB chunks typically finds redundancy within objects, such as when parts of text pages change. This performs even better and eliminates about 50% of the traffic. Finally, using 128-byte chunks can capture changes at the level of paragraphs in text, detecting about 55% redundancy. When we assume a cache that can store multiple days of traffic, the potential bandwidth savings further increase. As shown in Figure 3 (a), we could eliminate about 40% redundancy with 64 KB chunks (up from 30%), 60% with 1 KB chunks (up from 50%), and 65% with 128-byte chunks at the end of the week (up from 55%).

To determine the storage needed to achieve this redundancy, we simulate cache behavior of one week's traffic with an LRU cache replacement policy. Figure 3 (b), (c), and (d) show the byte hit ratio as a function of cache storage size in the United States, China, and Brazil, respectively. Missing values for 128-byte and 1 KB chunks in the figures are due to the limit of memory required for the simulation. As expected, the small chunk sizes yield much higher byte hit ratio than the large chunk sizes given the same cache storage size, and the byte hit ratio increases with larger cache storage. To achieve the ideal byte hit ratio, we need roughly 50 GB cache space for the United States, 500 GB for China, and 10 GB for Brazil. The difference in cache storage requirement is partly due to the volume of the traffic in these countries – while China's traffic is about 559 GB, Brazil's traffic is only about 44 GB.

This redundancy analysis result differs from earlier work [1] that analyzed data from smaller organizations with relatively homogeneous users. Our results use data covering diverse worldwide populations, and we show for the first time that there exists much redundancy in a large scale data set as well. The tradeoff, of course, is that much more storage is needed to capture this level of redundancy, but given the low price of storage relative to bandwidth, this tradeoff should be a simple choice in most countries.

3.4 Aborted Transfers

We find a significant portion of transfers are *cancelled* by clients before they fully fetch the objects. On any given day, roughly 17% of transfers are terminated early. Had these been allowed to complete, they would have accounted for 80% of the total traffic that would have resulted. However, with the early termination, they account for only 25% of the actual traffic generated. Some fraction of this traffic stems from ongoing transfers stopping when the user moves to a new Web page. However, most of the byte volume appears to be from previewing video files. Chunk-based caching solutions may be important here, since many of these files are never downloaded completely, making it hard to cache them in object-based caching solutions.

We next examine object size distributions. Figure 5 (a) presents the complementary cumulative distribution function (C-CDF) of object sizes. We see that while most of the cancelled objects ("Cancelled") are small, there exist very large objects ranging up to 2 GB. It is also interesting that most of the downloaded object sizes are less than 10 MB, regardless of being cancelled ("Cancelled-Download") or not ("Normal"). This implies that users typically tire of the content within the first 10 MB, something that may be useful to designers of video servers.

Figure 5 (c) and (d) compare content types of cancelled and normal objects in terms of the number of objects and total bytes. While text is the most frequent content type in both cancelled and normal objects, its percentage is extremely large in cancelled objects, exceeding 90%. In terms of total bytes served in terminated transfers, most come from a small fraction of video/audio/application objects.

Finally, Figure 5 (b) compares the distribution of cancelled transfers by region. We see that the first world countries cancel transfers more frequently than the developing world. The median cancel percentage difference is almost twice as large. One possible explanation for this is that clients from the developing countries are accustomed to slow transfers, and wait until the download finishes, to avoid wasting expensive bandwidth. On the other hand, clients from the OECD group cancel transfers more frequently since they have ample bandwidth, and cancelling does not negatively affect them.

4. RELATED WORK

There has been much research on network traffic analysis. Wolman et al. [26] analyze Web traffic of a university campus and enterprise network, and show that cooperative Web caching is most effective in smaller populations. Saroiu et al. [22] analyze client traffic at a large university to understand media download popularity. Anand et al. [1] focus on redundancy in university and enterprise network traffic, and propose that an end-to-end redundancy elimination solution can obtain most of the middlebox's bandwidth savings. In terms of the developing world context, Du et al. [8] analyze Web proxy access logs from Internet cafes and kiosks in Cambodia and Ghana. Anokwa et al. [2] focus on the high latency problem in the developing world, and propose a simple flow based classification solution. In terms of scale and granularity, our work covers a much larger number of users, a more diverse client population, and captures more of the

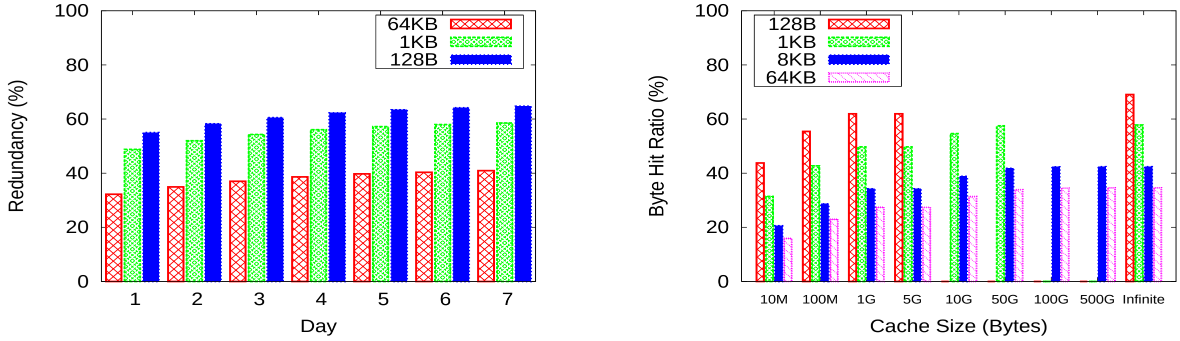

(a) Cumulative Redundancy (b) Byte Hit Ratio in the United States (213 GB)

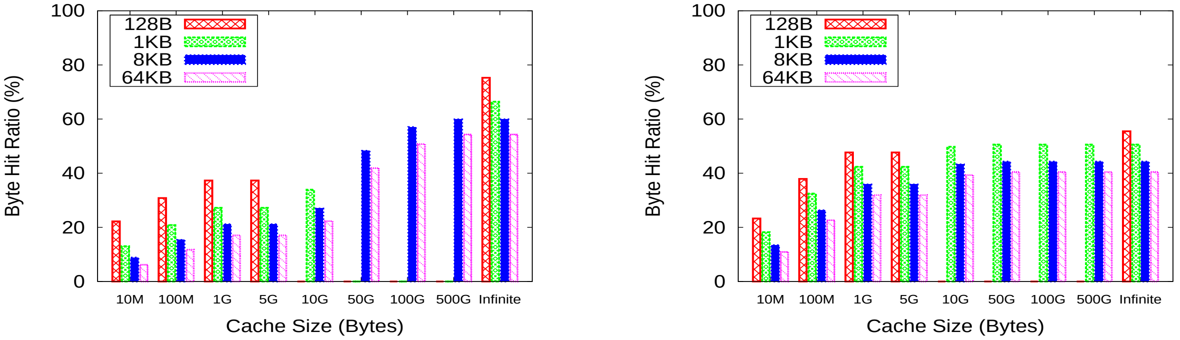

(c) Byte Hit Ratio in China (559 GB) (d) Byte Hit Ratio in Brazil (44 GB)

Figure 4: Redundancy and Caching

traffic than these previous efforts. While ipoque [10, 11] analyzes petabytes of user traffic from one million users worldwide, their main focus is peer-to-peer (P2P) traffic, not an analysis of Web content and caching.

We believe that our research runs complementary to two lines of work in developing areas – extending access and providing more local storage for caching. DakNet [19] and KioskNet [23] build Delay-Tolerant Networks (DTNs). WiLDNet [18] improves the performance of WiFi-based Long Distance (WiLD) networks in real-world deployments. RuralCafe [7] supports an efficient web search over intermittent networks by providing an expanded search query interface. HashCache [4] is a cache storage engine which is able to store a large number of objects with much less or no memory index, and Wanax [9] is a wide-area network (WAN) accelerator which could offer high performance with slow disks and limited memory.

5. CONCLUSIONS

In this work, we have used traces from a large, global proxy cache to help understand the difference between first-world and developing-world traffic. While some of our results, such as object size distributions, are in-line with previous research, we believe that our observations regarding download type distributions and the desire for rich media in the developing world are new contributions. We also find a great deal of redundancy in cache-miss traffic, which can be exploited over longer timeframes than existing Web caches support. We believe that this research provides a good case for increasing cache storage sizes, and augmenting object-based proxy caches with content-based chunk-oriented caching approaches. We consider this study a first glimpse at this traffic, and believe that a more detailed analysis is warranted.

6. ACKNOWLEDGMENT

We would like to thank the anonymous NSDR reviewers. This research was partially supported by NSF awards CNS-0615237 and CNS-0916204.

7. REFERENCES

[1] A. Anand, C. Muthukrishnan, A. Akella, and R. Ramjee. Redundancy in network traffic: findings and implications. In *SIGMETRICS '09*, pages 37–48, New York, NY, USA, 2009. ACM.

[2] Y. Anokwa, C. Dixon, G. Borriello, and T. Parikh. Optimizing high latency links in the developing world. In *WiNS-DR '08: Proceedings of the 2008 ACM workshop on Wireless networks and systems for developing regions*, pages 53–56, New York, NY, USA, 2008. ACM.

[3] A. Badam and V. S. Pai. Beating Netbooks into Servers: Making Some Computers More Equal Than Others. In *Proc. 3rd ACM Workshop on Networked Systems for Developing Regions (NSDR)*, Big Sky, MO, 2009.

[4] A. Badam, K. Park, V. Pai, and L. Peterson. Hashcache: Cache storage for the next billion. In *Proceedings of the 6th conference on Networked Systems Design and Implementation (NSDI'09)*, 2009.

[5] P. Barford and M. E. Crovella. Generating representative Web workloads for network and server performance evaluation. In *Proceedings of Performance '98/SIGMETRICS '98*, pages 151–160, July 1998. Software for Surge is available from Mark Crovella's home page.

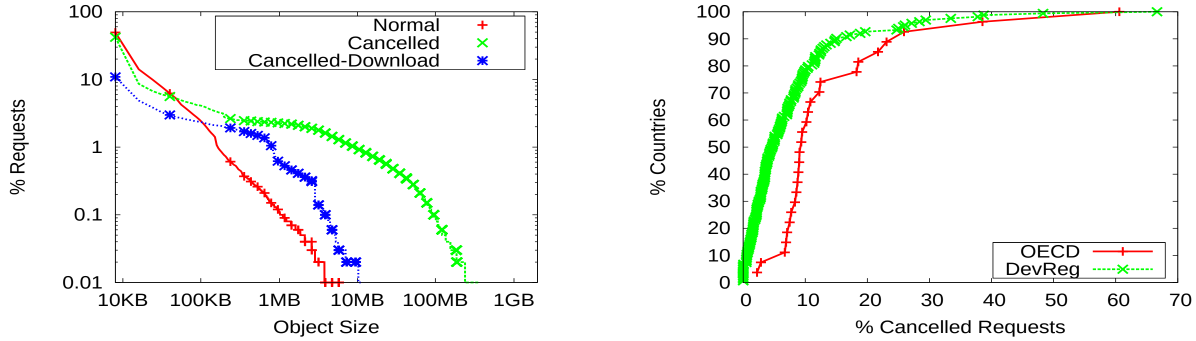

(a) Cancelled Object Size C-CDF (8KB Bin) (b) % Cancelled Object CDF

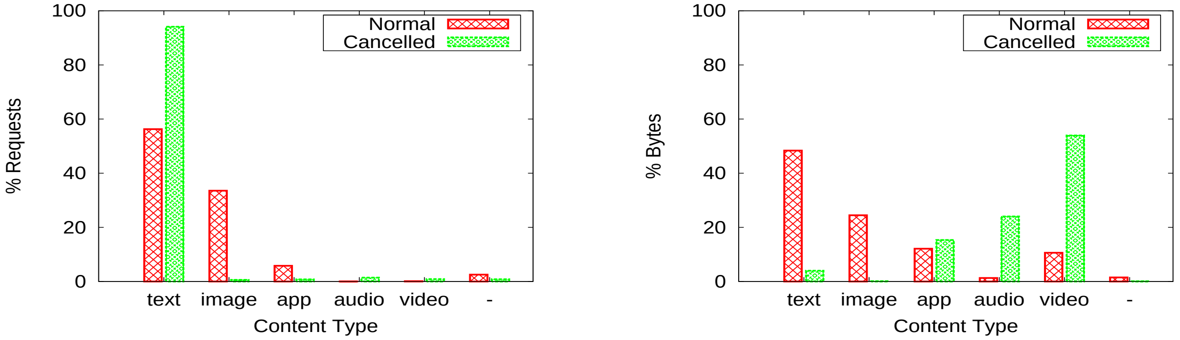

(c) Content Type By Requests (d) Content Type By Bytes

Figure 5: Cancelled Object

[6] A. Chankhunthod, P. B. Danzig, C. Neerdaels, M. F. Schwartz, and K. J. Worrell. A hierarchical internet object cache. In *USENIX ATC*, pages 153–163, 1996.

[7] J. Chen, L. Subramanian, and J. Li. Ruralcafe: web search in the rural developing world. In *WWW '09: Proceedings of the 18th international conference on World wide web*, pages 411–420, New York, NY, USA, 2009. ACM.

[8] B. Du, M. Demmer, and E. Brewer. Analysis of WWW traffic in Cambodia and Ghana. In *WWW*, 2006.

[9] S. Ihm, K. Park, and V. Pai. Wide-area network acceleration for the developing world. In *Proceedings of the USENIX Annual Technical Conference*, June 2010.

[10] ipoque. Internet Study 2007. `http://www.ipoque.com/resources/internet-studies/internet-study-2007`.

[11] ipoque. Internet Study 2008/2009. `http://www.ipoque.com/resources/internet-studies/internet-study-2008_2009`.

[12] B. A. Mah. An empirical model of http network traffic. In *INFOCOM '97: Proceedings of the INFOCOM '97. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Driving the Information Revolution*, page 592, Washington, DC, USA, 1997. IEEE Computer Society.

[13] C. Maltzahn, K. J. Richardson, and D. Grunwald. Performance issues of enterprise level web proxies. In *In Proceedings of the SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, 1997.

[14] MaxMind. `http://www.maxmind.com/`.

[15] J. Mogul, B. Krishnamurthy, F. Douglis, A. Feldmann, Y. Goland, A. van Hoff, and D. Hellerstein. Delta encoding in HTTP. RFC 3229, January 2002.

[16] J. C. Mogul, Y. M. Chan, and T. Kelly. Design, implementation, and evaluation of duplicate transfer detection in HTTP. In *USENIX NSDI*, 2004.

[17] Organisation for economic co-operation and development. `http://www.oecd.org/`.

[18] R. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. Wildnet: Design and implementation of high performance wifi based long distance networks. In *USENIX NSDI*, 2007.

[19] A. S. Pentland, R. Fletcher, and A. Hasson. Daknet: Rethinking connectivity in developing nations. *Computer*, 37(1):78–83, 2004.

[20] PlanetLab. An open testbed for developing, deploying and accessing planetary-scale services, September 2002.

[21] Riverbed Technology, Inc. `http://www.riverbed.com/`.

[22] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy. An analysis of internet content delivery systems. In *OSDI '02: Proceedings of the 5th symposium on Operating systems design and implementation*, pages 315–327, New York, NY, USA, 2002. ACM.

[23] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost communication for rural internet kiosks using mechanical backhaul. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, pages 334–345, New York, NY, USA, 2006. ACM.

[24] N. T. Spring and D. Wetherall. A protocol-independent technique for eliminating redundant network traffic. In *ACM SIGCOMM*, 2000.

[25] L. Wang, K. Park, R. Pang, V. Pai, and L. Peterson. Reliability and security in the CoDeeN content distribution network. In *Proceedings of the USENIX Annual Technical Conference*, 2004.

[26] A. Wolman, G. M. Voelker, N. Sharma, N. Cardwell, A. R. Karlin, and H. M. Levy. On the scale and performance of cooperative web proxy caching. In *Symposium on Operating Systems Principles*, 1999.