WEB TRAFFIC ACCELERATION OVER CELLULAR NETWORKS BY USING A COMPRESSING PROXY SERVER

Yair Toaff and Ariel J. Frank

Department of Computer Science, Bar-Ilan University, Israel

Keywords: Internet, cellular, proxy, compression, Performance Enhancing Proxy (PEP), GSM, HSCSD

Abstract: Many Web clients today are connected to the Internet via low speed computer links such as cellular connections. In order to efficiently use the cellular connection for Web access, the connection must be accelerated using a Performance Enhancing Proxy (PEP) as a gateway to the Web. In this paper we investigate the challenges created by the use of PEP. In order to mitigate the performance bottlenecks, we enhanced the PEP by adding to it a "pre-getting" ability. We tested the enhanced PEP over GSM and HSCSD cellular networks, as well as using a cellular network software simulator. Our experiments with enhanced PEP achieved the following results: 1) Average improvement of about 60% over the HSCSD network throughput can be achieved, 2) The Web page structure has a significant impact on the resulting performance. In conclusion, using an enhanced PEP can make the experience of browsing the Web over cellular networks much faster and less frustrating.

1 INTRODUCTION

Cellular access to the Internet began about a decade ago. The general characteristics of data transmission over cellular networks are a high bit error rate when compared to regular phone lines, frequent disconnects, and a generally slow transfer rate when compared to other means of data transmission. For these reasons, effective Internet browsing over cellular networks requires some means of acceleration. One such mean is the use of a compressing proxy server. The compressing proxy server is the focus of this paper.

We will first review the two types of technical challenges that we wish to analyze and resolve.

The first challenge is related to Internet browsing over cellular networks in general. The component issues are the time required to create a TCP connection (RFC1072 1988, Stevens 1996), and the synchronous protocol that the browser implements (Toaff 2002). Looking at the graph of the bandwidth used while downloading a Web page over a GSM (Global System for Mobile Communications) connection (Scourias 1996), we can see (Figure 1) that the bandwidth is not fully utilized.

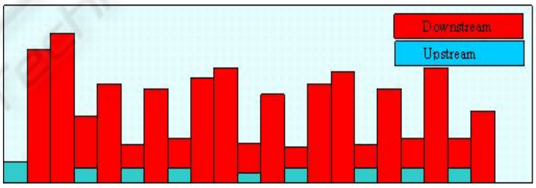

Figure 1: Web page download bandwidth graph

A partial solution for both these challenges is the use of a Performance Enhancing Proxy (PEP). PEP servers can usually provide a solution for the low bandwidth of cellular networks by using compression. Both client and server compress the data transferred over the cellular link. The data is then uncompressed at the target (Liljeberg 1996, Brown and Singh 1997, Toaff 2002).

The second set of issues is related to the use of PEP. When using PEP, the average actual bandwidth used drops significantly. It would be expected that if the compression ratio is about 1:3 – the resulting time ratio would be similar. What we have found, however, is that the average actual bandwidth used drops but the time ratio obtained is only about 3:4.

This phenomenon is the result of several reasons. The time needed for a request to reach the server remains the same but compression times must be added, causing a delay. There is less data to transfer

Yamada H. and Kawaguchi A. (2004).
A METHOD FOR THE PERFORMANCE ANALYSIS OF INTEGRATED APPLICATION SERVICES - Simulating the execution of integrated application services coordinated by workflow-driven broker-servers.
In *Proceedings of the First International Conference on E-Business and Telecommunication Networks*, pages 40-45
DOI: 10.5220/0001382900400045
Copyright © SciTePress

so the periods when data is downloaded are shorter. In the experiments described in section 5.1, although the size of the downloaded data is about 30% of the original data (120KB instead of 379KB), the download time drops only to about 50% of the original time (160s instead of 345s). This result still indicates inefficient use of the available bandwidth. In the advanced cellular networks, such as HSCSD (High Speed Circuit Switch Data) (ETSI 1997), the problem appears to be more severe (data ratio about 1:3, time ratio about 2:3) due to the asymmetric nature of the network (Toaff 2002).

The paper is organized as follows:

- Section 2 reviews related research works.
- Section 3 introduces the logical model for the pre-getting PEP.
- Section 4 describes the environment used for the experiments (both software and hardware).
- Section 5 describes the experiments that were carried out and the conclusions reached from them.
- Section 6 finally concludes with a discussion and suggests additional issues that may require further investigation.

2 RELATED RESEARCH WORKS

In this section we briefly review previous research works in related topics.

These related topics are:

- Data communication over cellular networks.
- Internet browsing over cellular networks.

2.1 Data communication over cellular networks

Several research works (Xylomenos and Polyzos 1999, Brown and Singh 1997, Vardalis 2001, Balakrishnan 1997, Tsaoussidis and Matta 2002) discuss the problems with data communication over cellular networks and suggest several methods of improving the existing TCP protocol. An important idea mentioned in these papers is data compression over the cellular connection but the idea itself was not implemented and tested, and the challenges caused by the use of compression were not discussed.

2.2 Internet browsing over cellular networks

Article (Liljeberg 1996) suggests a client-server based proxy that handles the pre-getting of page embedded objects. This idea is similar to our work; however, content compression was not implemented. The paper simply states that the results of compression are easily predicted. But as implied above, the results may not be as clear-cut as that paper predicts.
The main problem we have found with the described work is that it ignores new asymmetric cellular networks architectures (such as HSCSD) and assumes a network supplying stable bandwidth and equal bit rate for both upstream and downstream data transfer. In the HSCSD network we used, the downstream bit rate is 3 times the upstream bit rate. This influences performance since the time required for the request to reach the server remains the same, but the time required for the response to reach the client decreases by a factor of 3 (in HSCSD compared to GSM). This means that the full utilization of the downstream channel will require that 3 times the number of requests be sent for the same bandwidth. This network asymmetry causes the aforesaid problems to become more significant. Due to these problems, pre-getting is needed to enhance the utilization of the downstream channel and compression must be applied to both channels.

3 LOGICAL MODEL FOR ENHANCED PEP

The main object of our research has been the enhancement of an existing PEP with the purpose of resolving the problems previously encountered. We have found that adding a pre-getting mechanism to PEP achieves this goal.
In the pre-getting method, no speculative algorithms are required. Instead, pre-getting evaluates and modifies the way the browser requests information from Web servers and how the browser validates that information.
The pre-getting method works as follows:

1. The proxy client parses the HTML page, obtains a list of the embedded objects, compresses it and transmits the list of required objects to the server.
2. The proxy server builds the HTTP requests from the list, retrieves and compresses them, and returns them to the client.
3. The proxy client decompresses the responses and caches them.

4. When the browser requests the elements, they are ready for delivery, stored in the proxy client's cache.

The result of using pre-getting is that all the embedded objects from a certain HTML page are downloaded without time gaps. The idle download time previously observed is eliminated. The observable result is that the browser finishes downloading the page faster.

The parameter we have used for measuring the resulting improvement is the time required to download a Web page. We assume that the user is willing to "pay" for the reduced download time with some loss in the quality of the pictures contained in the page.

As the compression ratio can be assumed to be constant, therefore the improvement can be measured by the time required for the Web page to download. Our goal in enhancing the PEP is for the achieved time ratio to be similar to that of the compression ratio.

4 EXPERIMENT ENVIRONMENT

This section describes the environment we have used for our experiments. The description includes the software used in the experiments, the various components, and the hardware infrastructure.

4.1 Software

The general characteristics of the system are:

- Use of a client/server system.
- The client is connected to the server by a cellular connection.
- Both client and server support compression.

The commercial software used was the Apache Web Server and MS Internet Explorer browser.

Client

The regular compressing proxy client is composed of a communications module, a compression module, and a multiplexor module in charge of multiplexing the multiple browser connections into a single server connection.

To the enhanced compressing proxy client we added an HTML parser module and a cache module.

Server

The regular compressing proxy server is composed of a communications module, a compression module, and a multiplexer module. To the enhanced compressing proxy server we added a URL handler module for transforming the URLs into full HTTP requests.

4.2 Cellular network simulator

A software simulation of a cellular link was used due to the high cost of airtime on cellular networks and the requirement of the cellular link to supply constant bandwidth thus enabling the results to be compared.

4.3 Hardware infrastructure

Two MS Windows based workstations were used for the experiments. The communication hardware for the experiments on the real cellular network was a PCMCIA Nokia card phone. This card phone supports both GSM and HSCSD network connections.

For the simulator experiments, a 100Mb/s LAN was used to connect the computers. The experimental configuration is described in Figure 2.

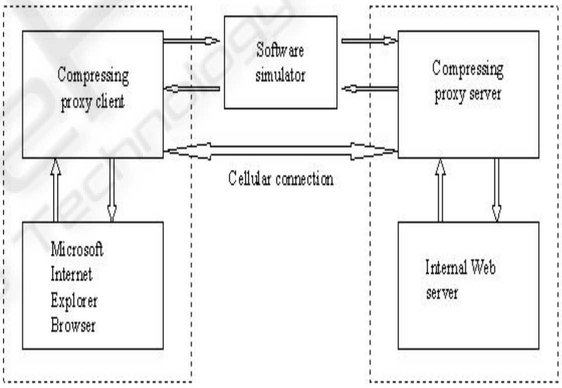

Figure 2: Experimental configuration overview

5 EXPERIMENTS AND CONCLUSIONS

This section describes the experiments. The first set of experiments analyzes performance over real cellular networks (GSM and HSCSD). The second set of experiments used the cellular network simulator. Following the description of each experiment, the resulting conclusions are discussed.

5.1 Real cellular network experiments

Our experiments using the pre-getting method were carried out as follows. A variety of Internet sites were copied onto a local Web server (to eliminate delays caused by communication with the content providers). Each of the chosen sites was downloaded three times:

- Without use of a proxy.
- Using the regular PEP.
- Using the pre-getting PEP.

For each download, 3 parameters were recorded:

1. The total amount of data.
2. The time required for download.
3. The average transfer rate.

After each download, the browser's caches were flushed (so the data and time totals could be compared).

Conclusion 1. The pre-getting method improves the compressing proxy bandwidth usage by up to 90% of the measured bandwidth with no compression used.

As can be seen in Table 1, the pre-getting method significantly improves the utilization of the available bandwidth.

Table 1: Cellular downstream average transfer rate (in bps)

Methods	Cellular transfer rate	
	GSM	HSCSD
Theoretical	9600	43200
No compression	8765	18595
Regular PEP	6062	9830
Pre-getting PEP	7946	17121

On the GSM network, the used bandwidth dropped by about 30% when using a regular PEP. The use of the pre-getting PEP lowered the drop to just under 10%.

On the HSCSD network, the used bandwidth dropped to about 50% when using a regular PEP. The use of pre-getting PEP lowered the drop also to just under 10%.

The improvement in the used bandwidth brings the resulting time ratio closer to the compression ratio. The 10% drop in the bandwidth utilization is due to the data passing more stages and to the computations requiring additional time (parsing, (de)compression).

5.2 Cellular network simulator experiments

The experiments of the pre-getting methods in the cellular network simulator were done in a similar manner, except that in this case the simulator replaces the cellular network

5.2.1 Imitating the cellular network

The first phase of the simulator experiments was to verify that the simulation closely simulates real network conditions, enabling reliable use of the results. This was tested using the parameters of the HSCSD network.

The parameter that was changed between the experiments is protocol overhead. The overheads that were tested included 148 bytes, 248 bytes and 348 bytes per frame.

Conclusion 2. The simulator imitates the behavior of the real cellular network with good precision.

It appears that simulated results based on a 23% loss in the bandwidth utilization (overhead of 348 bytes per frame), result in an effective bit rate that is almost identical to the HSCSD experiments.

5.2.2 Web page structure influence on the performance

Analyzing the structure of the Web pages used we identified that the number and size of embedded objects in a page affects the download time. In sites that contain a small number of pictures, (i.e., http://www.google.com – 2 small pictures), the time difference between using a regular PEP and a pre-getting PEP is very small.

On the other hand, there are sites that contain a large number of small pictures, (i.e., http://www.barnesandnoble.com – when used had 51 pictures, 39 of them with sizes below 1K). On these sites, the improvements obtained by utilizing a pre-getting method are very significant. The resulting download time dropped to about 30% of the download time using a regular PEP.

The average improvement on all the Web pages on a local server was about 50%. The bandwidth utilization ratio on the "small" sites was about 1.3 (i.e., the bandwidth used without PEP was 1.3 times higher than with pre-getting PEP), while on the "large" sites the ratio was about 0.7. The following 2 experiments analyzed the relationship between Web page structure and pre-getting performance, by using different Web pages built specially for this purpose.

Experiment 1 description

To further analyze the relationship between Web page structure and performance we added two additional pages that represent the extreme cases. The first page contains 6 pictures larger than 100KB and 8 pictures larger than 50KB. The second page contains 24 pictures smaller than 1KB and 28 pictures smaller than 5KB. We will refer to these 2 pages as big and small, respectively.

Conclusion 3. The performance of the pre-getting method depends on the Web page structure.

Table 2: Experiment results on "extreme" cases

		Big	Small
No Compression	Time (s)	306.6	83.6
	Total data	1252	163.3
	Average KB/s	32.6	15.6
Compressing proxy	Time (s)	158.7	64.9
	Total data	373.3	73.4
	Average KB/s	18.8	9.05
Pre-getting proxy	Time (s)	90.8	20.2
	Total data	373.3	73.4
	Average KB/s	32.9	29.1

The experiment results on these pages were as expected (see Table 2). As can be seen, the bandwidth utilized while using the pre-getting PEP remains stable around 30KB/s (on both types of pages). The bandwidth, while not using PEP at all, is affected by the page structure and drops by about 50% and while using the PEP it is about 60% of the bandwidth without a proxy.

<u>Experiment 2 description</u>

The previous experiment tested extreme cases. In order to analyze cases between the two extremes we performed an additional experiment. For this experiment, we collected a series of 8 pictures with sizes from 43 bytes to 80KB. For each of these pictures, we built a Web page that contained the picture 15 times.

The effect of the picture sizes on the bandwidth utilization can be seen in Figure 3. The basic conclusion that a large amount of small pictures degrades bandwidth utilization still holds. Bandwidth utilization is in direct relationship to the pictures sizes. The only exception is in the case that the picture is compressed extremely well (from 30,567B to less than 3,000B). Even in these cases, the bandwidth utilization obtained by using the pre-getting PEP is higher than the bandwidth utilization obtained by using a regular PEP. In normal cases (compression rate of 50% to 80%), we have found that the bandwidth utilization is similar to bandwidth utilization without a PEP and even exceeds it.

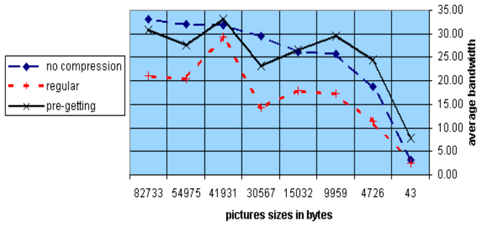

Figure 3: Bandwidth utilization in relation to pictures sizes

Only in the most extreme case (the smallest pictures) a drop in the used bandwidth was identified. But even with this drop, bandwidth utilization obtained by using other methods is exceeded.

Conclusion 4. A pre-getting compressing proxy gives us the ability to download large Web pages.

When attempting to download a large Web page using the cellular network a timeout on the browser was obtained. The browser attempts to download the page for 5 minutes and if the download is not completed the browser returns an error message. If we calculate the page size that will result in a timeout, we get 1.5MB (43,200/8 * 60 * 5 = 1,620,000) on a HSCSD network.

If we assume a compression ratio of 1:4 and bandwidth utilization that drops to about half, we achieve the ability to download Web pages that are twice as large without getting a timeout. When we use the pre-getting PEP (so the bandwidth used is about the same as without a proxy) we can download Web pages that are four times as large (i.e., up to 6MB in size).

6 DISCUSSION

The work presented here investigates the problems related to use of a compressing proxy server over cellular networks. The main problem was the effectiveness of bandwidth use. We saw the drop in bandwidth while using the PEP and studied its sources.

The next step was to offer a solution to this problem by adding a pre-getting mechanism to the compressing proxy server. We implemented this feature into the PEP and tested the effectiveness of the solution.

We saw that the pre-getting method boosted the performance of the PEP. With an average compression ratio of about 65% we obtained an improvement of about 60% in the download time, compared to 25% while using the regular PEP.

When we reviewed the results, we saw that the bandwidth ratio while using the pre-getting proxy server varied from 0.5 to 1.5. In order to identify the source of this inconsistency in a sterile environment, we built a cellular network software simulator.

Our experiments indicated that the structure of the Web page was the root cause of the variance in bandwidth utilization. We discovered that the best performance is on pages with a large number of embedded small pictures and that the worst performance is with pages holding a small number of large pictures. We constructed several sample pages to analyze the relationship between the sizes of the pictures embedded in the Web page to the pre-getting performance.

In conclusion, we have identified the problems which occur while using a regular compressing PEP, and introduced a mechanism that enhances an existing PEP in a way that enables improved Web surfing over a cellular network. Without this enhancement, Web surfing is slow, there is no direct relationship between quality loss and improvements in download time and it is practically impossible to download large Web pages. With this enhancement, the ratio between the loss in quality (the compression ratio) and the gain in time is close. In other words, if the user is willing to lose 90% of the data by reducing the quality of the pictures in the Web page, the page will be gotten 10 times faster.

Future research topics concerning the pre-getting method are:

- Adding compression for various data types – in this work, only GIF and JPG graphic files were explicitly compressed and at a constant ratio. All other types of files were compressed with the standard zlib compression.

- Improving the communication protocol – The communication protocol is implemented over TCP/IP. It can be improved by using the protocol improvements referred to in subsection 2.1.

REFERENCES

Balakrishnan, H., Padmanabhan, V. N., Seshan, S., and Katz, R. H., 1997. "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links", IEEE\ACM Transactions on Networking, 5(6), 756-769.

Brown, K., and Singh, S., 1997. "M-TCP: TCP for Mobile Cellular Networks", ACM Computer Communication Review, 27(5), 19-43.

ETSI 1997. "High Speed Circuit Switch Data (HSCSD) – stage 1", European Telecommunications Standards Institute, GSM specifications 02.34.

Liljeberg, M., Helin, H., Kojo, M., and Raatikainen, K., 1996. "Enhanced Services for World-Wide Web in Mobile WAN Environment", CS Technical Report No. C-1996-28, University of Helsinki.

RFC1072, 1988. "TCP Extensions for Long-delay Paths", http://www.cis.ohio-state.edu/Services/rfc/rfc-text/rfc1072.txt

Scourias, J., 1996. "Overview of the Global System for Mobile Communications", http://ccnga.uwaterloo.ca/~jscouria/GSM/gsmreport.html.

Stevens, W. R., 1996. "TCP/IP Illustrated", Vol. 1, Addison-Wesley, Chapters 18, 19, 21.

Toaff, Y., 2002. "Web Traffic Acceleration over Cellular Networks by using a Compressing Proxy Server", Master Thesis, Dept. of Computer Science, Bar-Ilan University, Ramat-Gan.

Tsaoussidis, V., Matta, I., 2002. "Open issues on TCP for Mobile Computing", Journal of Wireless Communications and Mobile Computing, Wiley Academic Publishers, 2(2), http://www.cs.bu.edu/techreports/pdf/2001-013-open-issues-tcp-wireless.pdf.

Vardalis, D., 2001. "On the Efficiency and Fairness of TCP over Wired/Wireless networks", Master Thesis, SUNY at Stony Brook, New York.

Xylomenos, G., and Polyzos, G., 1999. "TCP and UDP Performance over a Wireless LAN", Proc. IEEE INFOCOM.