x-ICP: A Mobile Web Caching Protocol

Wenzheng Gu and Abdelsalam (Sumi) Helal

Computer and Information Science and Engineering Department
University of Florida, Gainesville, FL 32611, USA
{wgu,helal}@cise.ufl.edu

ABSTRACT
Internet access by nomadic users roaming on a mobile network presents a scenario of access that is substantially different from the wired network. With current Web caching technology, accessing the Web while mobile is slow not only because of bandwidth limitations, but also due to moving away from the home Web caching proxy. Furthermore, current Web caching strategies ignore Web access patterns during periods of mobility. Studies of such scenarios are scarce. In this paper, we propose an Extended Internet Caching Protocol (x-ICP), in which the proxy server of the nomadic user's newly visited network can retrieve Web objects from its home network proxy server. Such a scheme decreases the response time for the requests by fetching an object from a usually nearby home network rather than from the origin site. We use trace-based analysis and analytic modeling methods to evaluate x-ICP. We draw several conclusions, all suggesting that x-ICP would be an effective Web caching approach for nomadic users.

Keywords
Internet Caching Protocol, Mobile Web Caching, Mobile IP.

1. INTRODUCTION
Web caching has been generally recognized as an effective approach to solve the long round-trip propagation delays, brought by the growth of the Internet [DAV99,BAR00,WAN99].

In a mobile network, the movement of nomadic users presents significant technical challenges to providing efficient access to the Internet. For a given individual mobile user, location and network point of attachment will vary in time. In this mobile environment proxy caching servers don't work as efficient as they do in a home network.

In this paper we introduce a new caching model and mobility extensions to the Internet Caching Protocol (ICP), to benefit the nomadic users. We call our protocol x-ICP. The main goal of x-ICP is to deal with the decreased cache hit rate on the proxy server when a user moves to a new location and attaches from a new network, which usually results in a longer response time. The x-ICP is based on the mobile IP protocol, which has been widely discussed as a network solution to mobility. Under Mobile IP, users don't have to modify their IP stack configurations manually each time they move into a new network [PER98,CAM00,CEL98,PER98a,VAL99]. But the design of Mobile IP makes the current caching system almost useless. X-ICP provides an opportunity to combat this problem by allowing the proxy server of the nomadic user's newly attached network (Foreign Proxy) to fetch a copy from a usually nearby home proxy server (Home Proxy). Figure 1 depicts this idea.

To gain more insight into the efficiency of x-ICP, we conducted some analysis by examining the traces of the request access logs collected by a proxy server at the CISE department of University of Florida, which identified some characteristics of collective user access patterns that were transformed into the foundation of x-ICP.

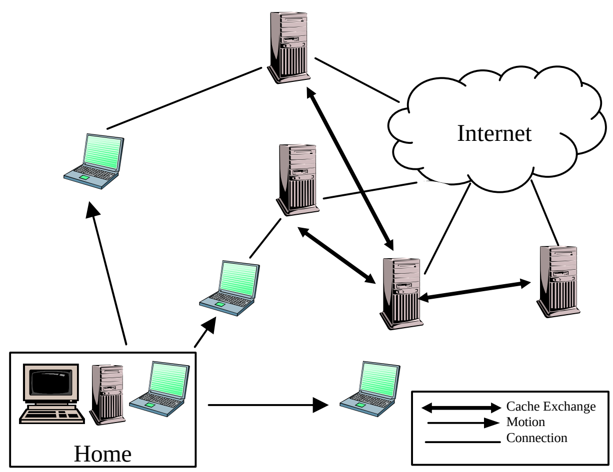

Figure1. X-ICP is deployed on the proxy servers belonging to the different networks. When PDAs or laptops are away from home, they can download copies of Web objects from Home Proxy

The remainder of this paper is organized as follows. Section 2 describes x-ICP in detail. Section 3 describes our traces and presents data from the traces. Section 4 develops an analytic model of the x-ICP and uses the model to study its performance. Section 5 summarizes the paper and provides conclusions.

2. OVERVIEW OF X-ICP

2.1 Fundamentals
First of all, a fundamental question is why do we extend ICP in particular?, why not extend other inter-caching protocols? Today, there are three main existing approaches on Web Caching, including broadcast probe, exemplified by Harvest and Squid, hash-partitioning of the object namespace among caches, like CARP and a directory service, first proposed in CRISP[TEW99,RAB98]. The later two infrastructures require pre-defined domains. But the uncertainty feature of nomadic users makes it hard to predict which proxy servers to be grouped together at the proxy setup time. Therefore, using short ICP messages to exchange caching information between Foreign Proxy and Home Proxy could be an effective way.

Several improvements have been made on ICP toward x-ICP. For instance, current shared cache architectures face a dilemma[TEW99]. On one hand, people wish to group as many proxy servers among a large number of clients to achieve good hit rates. On the other hand, as more proxy servers are involved in to service more clients, the response time they provide to any one client worsens due to the complex hierarchy. In x-ICP, such a problem doesn't exist. There are no siblings and parent. The Foreign Proxy will only contact with the Home Proxy where a great number of objects were previously downloaded by the same user. This will result in a relatively high hit rate compared to visit other proxy servers. Moreover, because of no broadcasting, the Internet traffic is reduced and response time is shortened because of no need to wait for the slowest response from all probed nodes.

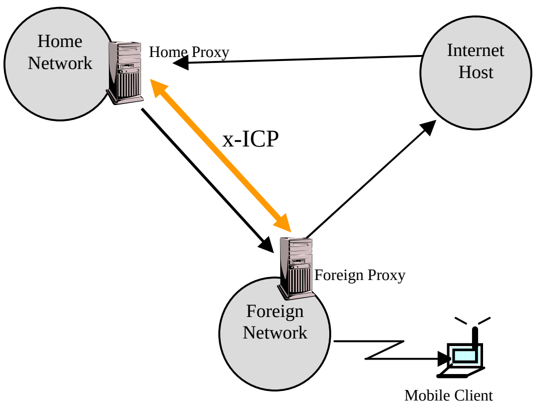

Figure 2. Applying x-ICP to Mobile IP's triangle routing

X--ICP is based on Mobile IP. Although Mobile IP brings significant conveniences for the nomadic users, there are still some unavoidable problems with it. One of them is that the triangle routing completely bypasses the Home Proxy server where a lot of Web objects are cached during previous Web accesses. Certainly, the Foreign Proxy mentioned here are those very close to the users, like the department proxy in a university, a corporation, or an ISP proxy. Compared to the higher level of proxy in a caching hierarchy, these proxy servers store many objects concentrated in a small area to cater to a small community. In the mobile IP infrastructure, when a nomadic node moves to a new location, he/she will send out an HTTP request, which is only checked by Foreign Proxy. If there does not exist any cached copy locally, the request will be then forwarded to the Web server at the far end of the Internet. But it's very likely a copy of the requested object is residing on the Home Proxy, which is usually nearby. The single-arrow line in Figure 2 shows the triangle routing of mobile IP and the solid double-arrow line indicates how we can take advantage of those copies closer to us. The key is when there is a cache miss happening on the Foreign Proxy, the request will be forwarded to Home Proxy server, if still a miss then to the origin site.

2.2 Mobile IP vs. X-ICP

Analogous to Mobile IP[PER98a], the main function of x-ICP can be divided into four major parts as illustrated in Table 1. They are Proxy Service Discovery, Sibling Proxy Server Registraten, Web Page Delivery, and Cache Duplication.

3. DESIGN DETAILS OF X-ICP

Before describing the detail procedures, we first need to introduce two processes, the Node Monitor and the Cache Copier, shown in Figure 3. The Nodes Monitor Module is responsible for maintaining several lists. One is the foreign user list. Those users who temporarily reside on this network are on this list. Another is the object list. Each user has a corresponding object list. All the objects' URLs, which are either cached locally, or fetched from origin servers, are on that list. Having detected the user's departure, the monitor sends over the URLs to the Cache Copier module on the user's Home Proxy.

The Cache Copier then simply checks whether they are stored locally, the expiration time and the expense to get it remotely. Then decides whether to download the object copies or not.

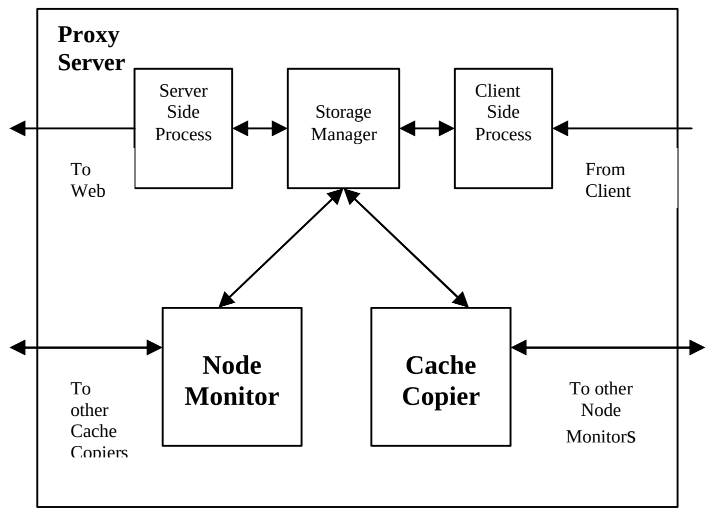

Figure 3. Modules on proxy server with x-ICP being deployed

3.1 Proxy Service Discovery

The proxy service discovery is based on the Service Location Protocol (SLP) [GUT99], which is a new IETF standards-track protocol designed to simplify the discovery and use of network resources such as Web servers, printers and etc.

3.2 Sibling Proxy Registration

Dynamic configuration is used in x-ICP. Regarding ICP, it is a lightweight message format used for communicating among Web caches. Caches exchange ICP queries and replies to gather information for use in selecting the most appropriate location from which to retrieve an object. Most of the implementations of current ICP version 2, like Squid, only allow configurations to be set at the installation time. When the server is running, no changes are allowed. In x-ICP, it's the nomadic node's responsibility to initialize the configuration updates. Therefore, the dynamic configuration is necessary.

3.2.1 Add Sibling Proxy Entries

Standard ICP is modified to adapt to our sibling proxy registration process. In x-ICP, when a nomadic node moves to a new location and finds out a new proxy server, the configuration file of the Foreign Proxy is modified and the access list of the node's Home Proxy is updated.

Table 1. Comparison of Mobile IP and x-ICP

	Mobile IP	X-ICP
Discovery	Home agents and foreign agents may advertise their availability on each link for which they provide service. A newly arrived nomadic node can send a solicitation on the link to learn if any prospective agents are present.	Proxy server agents advertise their services on their own links. A newly arrived nomadic node can retrieve service information from either the Service Agents or Directory Agents by Service Location Protocol (SLP).
Registration	When the nomadic node is away from home, it registers its care of address with its home agent either directly or through a foreign agent, which forwards the registration to the home agent.	When the nomadic node arrives at the new location, the configuration file of the Foreign Proxy server is modified to group the Foreign Proxy server with its Home Proxy server.
Delivery	In order for datagrams to be delivered to the nomadic node when it is away from home, the home agent has to tunnel the datagrams to the care-of address.	When a cache miss happens locally, the Foreign Proxy server forwards the request to the corresponding Home Proxy server first. If it is still a cache miss, then to the origin server.
Duplication	N/A	It is the Foreign Proxy's responsibility to monitor the status of each nomadic node. When a node is left, the Foreign Proxy needs to send the visited URLs related to that node back to its Home Proxy. The Home Proxy then decides which one to fetch.

Given the scenario in Figure 4, we will show how the registration works. A nomadic node was originally connected with "Net1", and later attached to the "Net2". The "Proxy4" is considered as its Home Proxy and the "ProxyE" is the Foreign Proxy. Having got to the "Net2", Node M inserts the following line into the configuration file of "ProxyE" as ICP required:

cache_host Proxy4.Net1 sibling HTTP-port icp-port. "cache_host" is the directive. Proxy4.Net1 is the host name or IP address of Home Proxy server. And the word "sibling" indicates the relationship between "Proxy4" and "ProxyE". There are no other relationships, like parent or multicast in x-ICP. At the mean time, "Proxy4" is notified by Node M to update its accessing list as following shows:

- acl src ProxyE ProxyE.Net2
- HTTP_access allow ProxyE
- icp_access allow ProxyE

This defines an access control entry named "ProxyE" which specifies a source IP address of the Net2 cache machine. The Net2 cache would then be allowed to make any request to both the HTTP and ICP ports. The above implementation is the same as the Squid implementation except the configuration file is modified by nomadic nodes dynamically.

Security needs to be carefully considered when the configuration is being updated. We are opening a door here to those malicious hackers if above implementations are used. For example, if someone puts an entry onto your list, saying look at its server as a sibling, then always satisfies the requests by returning irrelevant pages. This is dangerous. To prevent this from happening, we insert more information to the entry on the Foreign Proxy server. We also assume that the strong authentication mechanisms have been implemented by the mobile IP. Regarding the same case showed in Figure 4, we list an entry with one more constraint in bold font:

cache_host Proxy4.Net1 sibling HTTP-port icp-port ***care-of address***.

When Node M moves to Net2 and makes a request to register, the Foreign Proxy server also appends M's care-of address to the entry. Next time when an HTTP request comes in and if it's a cache miss, the proxy server looks up the caching list with the sender's care-of address. If there is a match, it forwards the requests to the corresponding Home Proxy server. Thus, no one can redirect other mobile nodes' requests. The security problem of updating Home Proxy's access list when the mobile node is away from home can be solved with the private key carried by the node.

3.2.2 Register with Node Monitor

One more thing that needs to be done is to register with Node Monitor. The Monitor maintains a Foreign Node Table on each proxy server. Having updated ACL, the monitor adds an entry to the table. The entry consists of three parts. The first is the mobile node's ID; the second is the host name or IP of its Home Proxy server. In the above case, the entry is like "M.Net1, Proxy4.Net1". And last, the Node Monitor creates a corresponding empty object list for the node for later object retrieval.

3.3 Web Objects Delivery

When a nomadic user makes a request via the Foreign Proxy, the Storage Manager of the proxy first checks the local disk. If the result is a cache hit, the object will be returned directly, and the URL of the object will be inserted into the object list of that user.
If it is a cache miss, the Foreign Proxy will forward the request to its sibling. And the proxy will fetch the copy if it is not out of date. Otherwise, the copy will be fetched from the origin site. In the last case, the URL of the object will be inserted into the user's object list, too.

3.4 Cache Duplication

The goal of Cache Duplication is to copy the cached objects of a nomadic user from Foreign Proxy to the Home Proxy when the user leaves. The Node Monitor on the Foreign Proxy and the Cache Copier on the Home Proxy are cooperated with each other to finish this task. When the Node Monitor detects the leaving of a nomadic node on the local network, it begins to negotiate with the Home Proxy of the node. It sends over the complete object list of that node to the Cache Copier of the Home Proxy. The Cache Copier asks the local Storage Manager to see whether some objects on the list are cached locally, and removes those URLs already cached. Thus the Cache Copier makes a shortened URL list and sends it back to the Node Monitor of Foreign Proxy. With the updated list, the monitor is able to send out the objects accordingly. And Cache Copier receives and stores them on the local disk of Home Proxy.

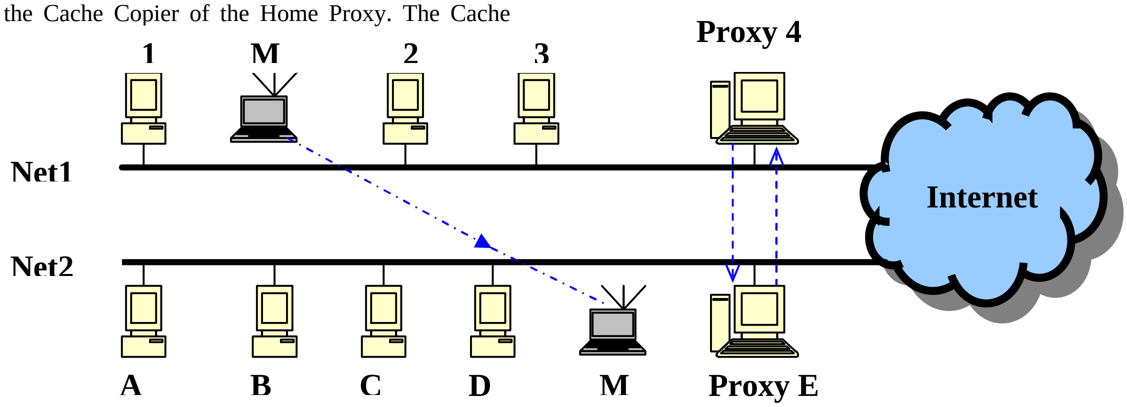

Figure 4. A nomadic node M moves from Net1 to Net2

Table 2. Average objects distribution per day on the different CISE sub-networks

	Sub-net 1	Sub-net 2	Sub-net 3	Sub-net 4	Sub-net 5	All others
Objects/day	420	2307	6960	7280	23535	426

Having done so, the Node Monitor also deletes the entry of the leaving nomadic node from the Foreign Node list. And the whole list is checked thoroughly. If there is no other nodes are sharing the Home Proxy in this entry, the corresponding “cache-host” line will be removed, too.

4. SIMULATION

4.1 Trace Collection and Characteristics

This and next sections pose and answer a number of questions about the potential performance of x-ICP. We will explore the bounds of x-ICP. Key questions include:

- What cache hit rates on the Home Proxy can be achieved?
- For what range of distances can x-ICP get an object from Home Proxy without introducing a significant delay?
- What is the speedup when x-ICP is deployed?

To answer these questions quantitatively, we have collected and analyzed Web access data from two environments First, our observations on the traces of user requests collected by a proxy server at CISE department of University of Florida during the period 10/29/2001 – 11/19/2001 are presented. On average there are 40,199 accesses everyday from more than 5 sub-networks. Table 2 shows the requests distribution from the different sub-networks and the corresponding population is in proportion to the number of requests. We selected three of them (sub-net 1,3,5) with different number of requests and ran the simulation separately to test the hit rate on the Home Proxy. Second, we also conducted experiments on a 19-day NLANR trace [WES01]. The data is collected on one of their ten proxy servers, named “startap” from Feb 24, 2002 to March 14, 2002. The feature of root node in a caching hierarchy made it less interesting for our analyses, but where we still could obtain comparable results.

4.2 Simulation methodology

The results presented in this section are based on Web caching simulations using above mentioned traces as input. We installed the Squid proxy server version 2.4 on two UNIX machines, and configure them as sibling. The proxy directly answering clients’ requests is considered as the Foreign Proxy and the other is the Home Proxy. In analyzing the cache behavior of these traces, we use the first three days of each trace to warm our Home Proxy caches before gathering statistics. Whether the object is cacheable or not is judged by the Squid proxy servers, and which objects are to be replaced are taken care of by Squid, too.

These traces have two primary limitations that affect our results. First, we assume the user linger time on the new location is one day long. That’s good enough for nomadic users. Second, each time when we start a new day simulation, the cache on the Foreign Proxy is set empty, because currently no research has been conducted on what kind of Foreign Proxy the nomadic user is likely to attach to. Therefore, the compulsory misses on the Foreign Proxy is a little higher than that in the reality.

To quantify the overall performance of the system and examine the impact of x-ICP, Figure 5 shows both the Home Proxy hit rates and Foreign Proxy hit rates in an about 20 days trace. The bars are labeled on the x-axis with days. The lower portion of each bar shows the hit rate that would be seen by a local Foreign Proxy acting on behalf of the nomadic user’s new attachment point. The upper portion of the bar shows the improvement in hit rate that would be seen by the nomadic user using x-ICP. Those cache hits are from user’s home proxy. In (a), for the Foreign Proxy bars, the average hit rate is 12.9% ranging from 0 to 29.1%. For the Home Proxy bars, the average hit rate is 22.1% ranging from 1.2% to 39.9%. In (b), for the Foreign Proxy bars, the average hit rate is 23.2% ranging from 9.1 to 72.6%. For the

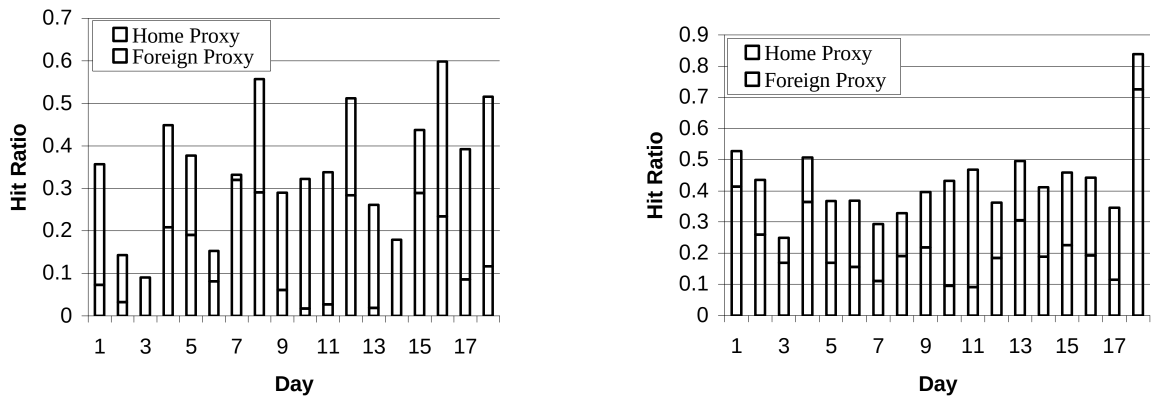

(a) Subnet 1 of CISE proxy log **(b) Subnet 3 of CISE proxy log**

Figure 5. Home Proxy and Foreign Proxy hit rates. Lower part of the bars are the local cache hit rates. The upper portions are the gain when we apply x-ICP to proxy servers. They are the hit rates on the Home Proxy when local misses happen on Foreign Proxy.

Home Proxy bars, the average hit rate is 19.8% ranging from 8.0% to 37.7%.

From the above data, the Foreign and Home Proxy hit rates changed when different networks were tested. Actually, behind that, the population and the number of requests of each sub-network play an important role. We briefly discuss this result here and the details can be found in [WOL99] to explain the relationship between the hit rate and population. Equation (1) summarizes the key formula from Wolman et al. yielding C_N, the aggregate object hit rate for a population N, where n denotes the number of requests. Thus, the hit rate is $H_N = p_c C_N$, where p_c is the probability that a requested document is cacheable. This model makes several assumptions: (i) object popularity follows a Zipf-like distribution, (ii) object popularity is uniform across the population, (iii) objects are cacheable with probability p_c independent of popularity, and (iv) inter-arrival time for object requests and updates are exponentially distributed with a parameter determined by the object's popularity. [GAD00] All these are fit for the nomadic environment so that we are able take advantage of it. In short, the bigger population the proxy serves, the higher the local hit rate is. According to [WAN99,GOL98], the hit rate at a given proxy is almost fixed, therefore, the higher the local hit rate is, the lower it is on the Home Proxy.

Since we are more interested in the behaviors of an individual or a group nomadic users of fairly small number, we believe that the data collected from the sub-net 1 and 3 are more likely to match our criteria. We eliminate the data from sub-net 5 because the requests from it take up half of the total requests of the whole department's, where there are hundreds of students and faculty members. The NLANR trace shows the smallest Home Proxy hit rate, which is about 12.2%, because it's one of the root cache of the national wide caching hierarchy system, which can be considered as serving a very large population and results in the very high local hit rate. In turn it reduces the Home Proxy hit rate. So, we didn't adopt its value, too.

Therefore, we calculate the average of the Home Proxy Hit Rates on sub-net 1 and 3, which is 21%. So, if x-ICP caching were possible, it would achieve a noticeable improvement in hit rate for nomadic users.

$$C_N = \int_1^n \frac{1}{Cx^a}\left(\frac{1}{1+\frac{mCx^a}{lN}}\right)dx, \quad C = \int_{1\le x\le n} \frac{1}{x^a}dx \qquad (1)$$

5. ANALYSIS

Briefly speaking here, having got the hit rate in section 4, we then calculate the boundary value of inter-proxy delay(RTT) in section 5.1. The enough hit rates on Home Proxy can compromise the inter-proxy delays. Certainly, the lower the average delay is, the higher the overall performance is.

With RTT, the potential distance (How far away a user can move from his Home Proxy) can be found out. Similar to Distance = Time * Speed, knowing the current network speed by some statistic data and the average delays from section 5.1, we calculate the potential inter-proxy physical distance in section 5.2.

5.1. Hit rate vs. Inter-proxy delay

In the previous section we considered the potential performance of x-ICP using our trace data. In this section, we will extend these results analytically to better understand the caching performance of x-ICP.

In the x-ICP scheme, a Foreign Proxy forwards a missing request to the nomadic user's Home Proxy to determine if (i) the Home Proxy holds the requested document, and (ii) that document can be returned faster than a request to server. Whether such cooperation is worthwhile will mainly depend on the inter-proxy latencies.

We begin by setting up a model to show the relationship between hit rate and inter-proxy delay based on the x-ICP infrastructure. Let's assume a group of people moves from the Home Proxy to the Foreign Proxy. Below are some defined terms:

- **R**: the total number of object requests made by the moving group people when they are attached to a new network.
- **C_h**: the hit rate on the Home Proxy when R requests were made.
- **C_f** : the hit rate on the Foreign Proxy when R requests were made..
- **D_h**: the average delay between the Foreign Proxy and the Home Proxy.
- **D_o**: the average delay between the Foreign Proxy and the origin server.

Therefore, we have the number of missed requests on the Foreign Proxy and forwarded to the Home Proxy, which is $R(1- C_f)$. Finally, we have

$$Speedup = \frac{R \cdot (1 - C_f) \cdot D_o}{R \cdot (1 - C_f) \cdot [D_h + (1 - C_h) \cdot D_o]} \quad (2)$$

The numerator denotes the total delay without x-ICP, and the denominator is the total time when x-ICP is deployed. Equation (2) can be reduced to equation (3).

$$Speedup = \frac{D_o}{D_h + (1 - C_h) \cdot D_o} \quad (3)$$

Lots of research has been conducted on the Internet round trip time (RTT) [WOR01,COT00]. We adopt the value 65ms for regional RTT within Europe and within North America, which is assigned to D_o. Under the assumption of an individual or a small number of nomadic users, we assign the value 21% from previous section to **C_h**.

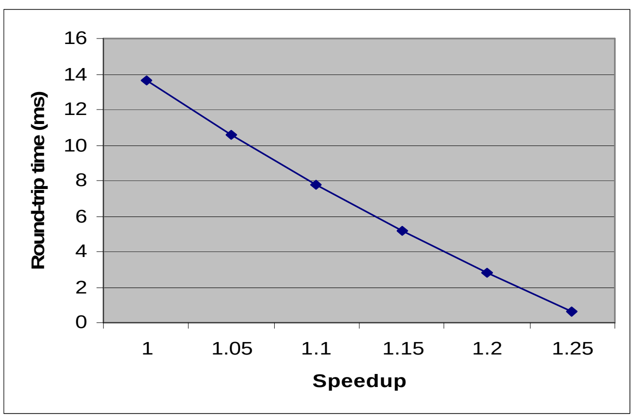

Figure 6. As x-ICP is deployed, downloading efficiency will be a function of inter-proxy roundtrip time according to equation (3).

In Figure 6, the speedup value illustrates how many times faster of the response time that people are expecting to achieve when using x-ICP than without it. As the graph shows, the higher the speedup, the less the Inter-proxy RTT is required. When a nomadic user attaches to a new network, the RTT between Home Proxy and Foreign Proxy cannot exceed 13.65 ms. Otherwise, people have to wait for longer time than they fetch the copies from the origin server directly. Furthermore, we can see that when people want to experience a 1.22 times shorter response time, the RTT cannot be more than 2ms.

5.2. RTT vs. Distance

The further study shows the relationship between RTT and distance. In [COT00], having measured 16 pairs of sites located in the U.S., Europe and Japan, Cottrell et al. were able to generate an equation (4), where Y indicates RTT and is the function of the distance (Kilometer), which is denoted by X in equation (4).

$$Y = 0.024X + 5.8887 \quad (4)$$

With some simple calculation, we conclude that the maximum distance x-ICP can support is about 321 kilometers (201 miles) when Y is 13.65 ms.

We have also collected some data on RTT in our lab. The simulation method in [ZHA93] was implemented. The result is shown in Figure 7. The seven hops are on the routing path, through which most Web objects requests from our lab are sent to the origin sites. The first six hops are on campus. The seventh hop is a carrier provider's node in another city. The definition of hop is the same as it is in traceroute algorithm. We tested the delays in the six different periods of a day, and observed that most delays are less that 2 ms in a campus wide network. According to Figure 6, the total response time can be 1.22 times faster when x-ICP is deployed on campus. Therefore, we are able to answer the questions asked at the beginning of the previous section.

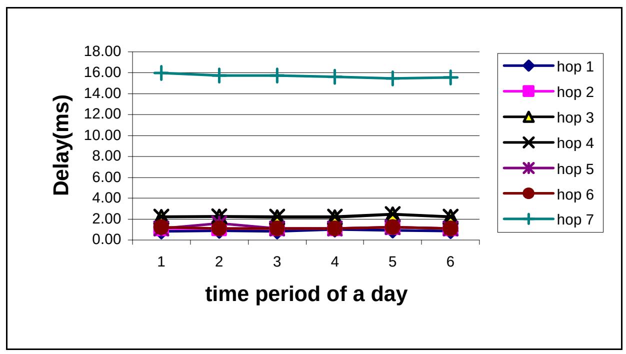

Figure 7. Round trip delays tested from a lab in CISE dept. of University of Florida. The first six hops are on the campus network, the seventh hop is on the ISP net.

- About 21% cache hit rates could be obtained on Home Proxy server.
- While moving, nomadic users can achieve as many as 1.22 times faster response time on the campus wide network.
- Nomadic users can attach to any new network without lowering the performance as long as the Foreign Proxy is not farther than 201 miles away from the Home Proxy.

In the above analysis, we used some empirical values. To further explore the x-ICP performance, we analyze the sensitivity of x-ICP on D_o and C_h. The D_o we used is the RTT measurement on a small network packet. When a Web object is downloaded, especially on a wireless network, the value of D_o is bigger. Figure

4-5 shows that the x-ICP speedup is not changed significantly when D_o grows. On the other hand, when the overall network speed becomes faster, the speedup drops if the delay between the two proxy servers is large. Figure 4-6 illustrates that to improve the performance of x-ICP, the hit rate on the Home Proxy is critical. When hit rate is bigger, speedup grows faster, for instance, when the hit rate is 50% on the Home Proxy, the speedup can be 2. When the hit rate on the Home Proxy is low, its impact on the speedup is not significant.

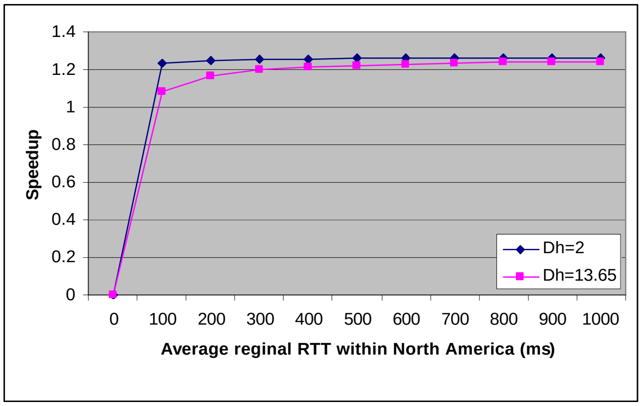

Figure 4-5. The impact of RTT values on speedup

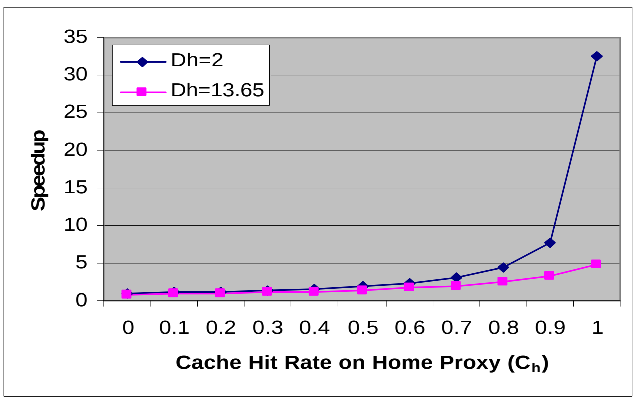

Figure 4-6. The impact of the cache hit rate of the Home Proxy on speedup

6. CONCLUSION AND FUTURE WORK

In this paper, we argued that Web caching is not benefiting nomadic users who change point of attachment to the Internet from time to time. We introduced x-ICP, a Web caching protocol for nomadic users who access the Internet through a mobile IP network. The design details of x-ICP are described, and some insightful statistics and analyses derived from a trace of a Web proxy server access logs are given. The results justify x-ICP in many ways. However, since the Internet users are moving faster and the devices are getting smaller. We will continue to work on the micro-mobility IP networks.

7. REFERENCES

[DAV99] Brian D. Davison. A Survey of Proxy Cache Evaluation Techniques. In Proceedings of the 4th International Web Caching Workshop, San Diego, California , March/April 1999. NLANR and CAIDA.

[PER98] Charles E. Perkins "Mobile Networking through Mobile IP," IEEE Internet Computing, January-February 1998.

[CAM00] Andrew T. Campbell "Design, Implementation, and Evaluation of Cellular IP," IEEE Personal Communications Magazine, August 2000.

[CEL98] Cellular IP HTTP://www.comet.columbia.-edu/cellularip/spec.htm

[PER98a] Charles E. Perkins Mobile IP Design Principles and Practices. Addison-Wesley 1998

[VAL99] Andras G. Valko Cellular IP - A New Approach to Internet Host Mobility, ACM Computer Communication Review, Jan 99.

[WAN99] Jia Wang A Survey of Web Caching Schemes for the Internet ACM Computer Comm.Rev., vol.29,no.5 Oct. 1999

[BAR00] Greg Barish and Katia Obraczka World Wide Web Caching: Trends and Techniques, IEEE Comm. Internet Technology Series, vol.38, no.5, May 2000, pp.178-184

[MOG00] J.C.Mogul Squeezing More Bits out of HTTP Caches, IEEE Network, vol. 14, no.3, May/June 2000, pp.6-14

[COF01] K. G. Coffman and A. M. Odlyzko Internet growth: Is there a “Moore’s Law” for data traffic? AT&T Labs – Research June 4, 2001

[CHI99] Cho-Yu Chiang, Ming T. Liu and Mervin E. Muller Caching Neighborhood Protocol: a Foundation for Building Dynamic Web Caching Hierarchies with Proxy Servers. In Proceedings of the International Conference on Parallel Processing, Fukushima, Japan , September 1999.

[BAL98] Hari Balakrishnan, Challenges to Reliable Data Transport over Heterogeneous Wireless Networks, Ph.D. Diss., UC Berkeley, 1998

[EFF93] Internet Statistics http://www.eff.org/Net-_culture/Statistics/internet_index_7-8-93.index

[TRA97] Traceroute White Paper http://www.inf-ormatik.unitrier.de/~smith/networks/tspec.html

[GOL98] Arthur Goldberg, Ilya Pevzner, and Robert Buff. Caching Characteristics of Internet and Intranet Web Proxy Traces. In Computer Measurement Group Conference, Anaheim, California, USA , December 1998.

[BOV02] C.J. Bovy, H.T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal and P. Van Mieghem Analysis of End-to-End Delay, Measurements in the Internet passive and active measurement workshop 2002

[WOL99] Alec Wolman, Geoffrey M. Voelker, Nitin Sharma, et al. On the scale and performance of cooperative Web proxy caching, 17th ACM Symposium on Operating System Principles, 1999

[GAD00] S. Gadde, J. Chase, and M. Rabinovich. Web Caching and Content Distribution: A View from the Interior. 5th Int. Web Caching and Content Delivery Workshop (WCW'00), May 2000.

[ZHA93] Y. Zhang and B. Bhargava, WANCE: A Wide Area Network Communication Emulation System, Proceedings of IEEE workshop on Advanced in Parallel and Distributed Systems, pp. 40-45, Princeton, NJ, Oct. 1993.

[WES01] Duane Wessels National Lab of Applied Network Research access logs

[WOR01] Worldcom Latency Statistics http://www1.worldcom.com/global/about/network/latency/

[COT00] L. Cottrell, W. Matthews and C. Logg Tutorial on Internet Monitoring & PingER at Stanford Linear Accelerator Center.

[RAB98] M. Rabinovich, J. Chase, and S. Gadde. Not all Hits are Created Equal: Cooperative Proxy Caching Over a Wide Area Network. In Proc. of the 3rd International WWW Caching Workshop, Manchester, England , June 1998.

[TEW99] R. Tewari, M. Dahlin, H. Vin, and J. Kay. Beyond Hierarchies: Design Consideratens for Distributed Caching on the Internet. In Proceedings of the 19th International Conference on Distributed Computing Systems, Austin, Texas , June 1999. IEEE.

[GUT99] Erik Guttman Service Location Protocol: Automatic Discovery of IP Network Services, In proceedings of the Internet Computing, July/August 1999 (Vol. 3, No. 4). IEEE