Evaluating Cooperative Web Caching for Emerging Network Technologies

Christoph Lindemann and Oliver P. Waldhorst

University of Dortmund
Department of Computer Science
August-Schmidt-Str. 12
44227 Dortmund, Germany
cl,ow@ls4.cs.uni-dortmund.de
http://www4.cs.uni-dortmund.de/~Lindemann/

Abstract While bandwidth for previous IP backbone networks deployed by Internet Service Providers typically has been limited to 34 Mbps, current and future IP networks provide bandwidth ranging from 155 Mbps to 2.4 Gbps. Thus, due to emerging network technologies cooperative Web caching among institutional proxy caches residing a wide area backbone network shows great promise to become an effective approach for reducing Web document access latencies in the global Internet. In this paper, we present a comprehensive performance study of four cooperative Web caching protocols. We consider the Internet cache protocol ICP, Cache Digests, the cache array routing protocol, CARP, and the Web cache coordination protocol, WCCP. The performance results are derived using trace-driven simulation. In particular, the effect of rapidly increasing bandwidth availability due to emerging network technologies on the performance achieved by these protocols is investigated. The presented results show that hash based protocols such as CARP and WCCP can effectively support cooperative web caching in high-speed backbone networks.[1]

1 Introduction

Cooperative Web caching means the sharing and coordination of cached Web documents among multiple communicating proxy caches in a backbone IP network or in an intranet. Cooperative caching of data has its roots in distributed file and virtual memory systems in a high-speed local area network environment. Cooperative caching has shown to substantially reduce latencies in such distributed computing systems because network transfer time is much smaller than disk access time to serve a miss. While bandwidth for previous IP backbone networks deployed by Internet Service Providers typically has been limited to 34 Mbps, current and future IP networks provide bandwidth ranging from 155 Mbps to 2.4 Gbps (see e.g., CA*Net-3 [1], G-WiN [4], or Internet-2 [5]). Thus, due to emerging network technologies cooperative Web caching among proxy caches residing a wide area backbone network shows great promise to become an effective approach for reducing Web document access latencies in the global Internet.

Protocols for cooperative Web caching can be categorized as message-based, directory-based, hash-based, or router-based. Wessels and Claffy introduced the Internet cache pro-

[1] This work was supported by the DFN-Verein with funds of the BMBF

tocol, ICP [13], which has been standardized and is widely used. As a message-based protocol, ICP supports communication between caching proxies using a simple query-response dialog. Directory-based protocols for cooperative Web caching enable caching proxies to exchange information about cached content. The information is compressed using arrays of bits, so called Bloom filters. Notable protocols of this class include Cache Digests by Rousskov and Wessels [10] and Summary Cache by Fan, Cao, Almeida, and Broder [3]. The most notable hash-based cooperative Web caching protocol constitutes the cache array routing protocol, CARP, introduced by Valloppillil and Ross [9], [12]. The rationale behind CARP constitutes load distribution by hash routing among Web proxy cache arrays. Cieslak, Foster, Tiwana, and Wilson introduced the Web cache coordination protocol, WCCP, as a router-based protocol. WCCP transparently distributes requests among a cache array [2].

To best of our knowledge, a comprehensive performance study of ICP, CARP, Cache Digests, and WCCP based on the same IP backbone network topology and workload characteristics has not been reported so far. Furthermore, the effect of rapidly increasing bandwidth availability to these cooperative Web caching protocols has not been investigated. We are also not aware of any performance study of WCCP v2.0. This paper tackles these issues. In Section 3, we present a comprehensive performance study for these cooperative Web caching protocols. The presented results are derived using a discrete-event simulator of a backbone network implemented using the simulation library CSIM [11]. We investigate cooperative Web caching from viewpoints of the client and the network. As performance measures, we consider user latency and bandwidth consumption. The goal of our study lies in understanding the behavior and limiting factors of cooperative Web caching protocols network for reducing Web document access latencies in the global Internet. We point out weaknesses of existing protocols for cooperative Web caching and state key issues for performance improvements. Particular, attention is paid to investigating the effect of rapidly increasing bandwidth availability due to emerging network technologies on the applicability of cooperative Web caching among institutional proxy caches residing a wide area backbone. We conclude from our performance studies that Fast backbone networks allow employment of hash routing protocols as CARP or WCCP for distributed caching in wide area networks. Furthermore, we conclude that emerging network technology has little impact on efficiency of distributed caching, because latency is dominated by download delays.

The paper is organized as follows. The simulation environment for evaluating the considered protocols for cooperative Web caching is introduced in Section 2. In Section 3, we present performance curves for the considered protocols for cooperative Web caching derived from the simulator and the measured trace data. Finally, concluding remarks are given.

2 Simulation Environment

In this section, we provide an overview over the simulation environment as well as the characteristics of the trace files, which have been used as input for our performance studies.

2.1 Network Model

In our experiments, we investigate the performance of cooperative caching for caching proxies distributed across the backbone network of a national Internet service provider (ISP). To simulate a particular network topology, we adopt the modeling approach for IP networks connecting clients and servers presented in [8]. They model the Internet as a hierarchy of ISPs. All clients (i.e., user agents) are connected to institutional networks, which are connected via regional networks to a national back- bone. National backbones are also connected. Caches are located at the access points between two networks to reduce the consumption of network bandwidth. As in [8], the simulator represents just one local and one remote national backbone network.

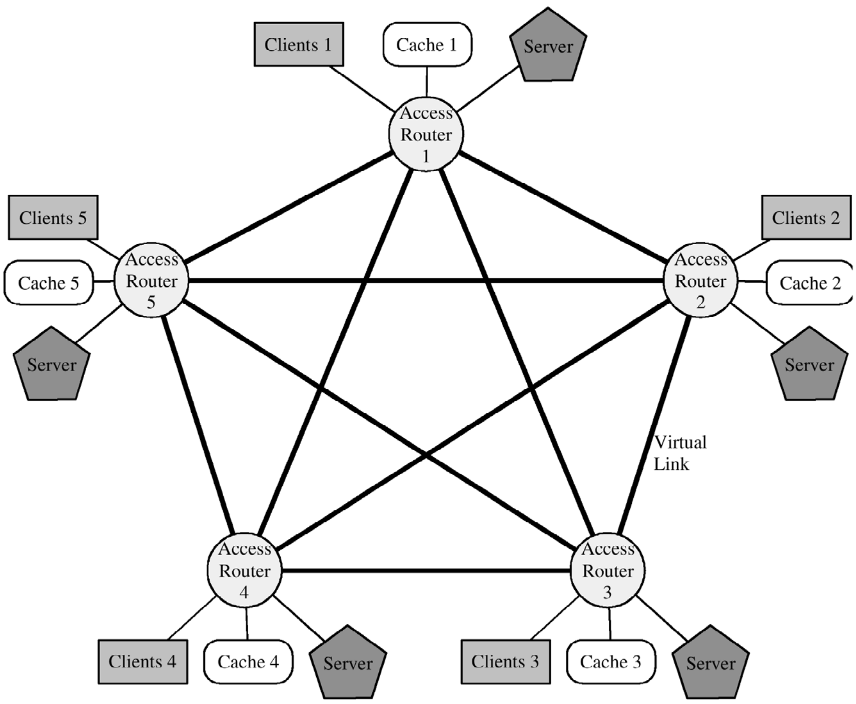

Figure 1. Considered network configuration with virtual links

Figure 1 shows the configuration of the local backbone network considered for the comparative performance study presented in Section 3. The configuration comprises of five caching proxies, each of them connected to an access router. The access routers are fully connected by virtual links. Clients and a local Web server are connected to each access router. The simulator assumes the network configuration to be fixed. Thus, we do not take into account failures and down times of any network component. The simulator comprises of the building blocks described above which can be utilized to assemble arbitrary configurations for backbone networks. These building blocks are CSIM simulation components modeling virtual links, access routers, clients, caches, and Web servers.

2.2 Characteristics of the Considered Workloads

To derive meaningful performance results for distributed Web caching in a national backbone network, we need to feed our simulator with representative workloads collected at top-level caches. Request streams for top-level caches show significantly different characteristics than request streams for institutional caches [6]. The performance studies presented in Section 3 are based trace files from log files of two different networks: (1) four top-level proxy caches of the German Gigabit research network, G-WiN [4] and (2) ten top-level proxy caches of the National Laboratory for Applied Network Research (NLANR) cache hierarchy [7]. At the time of trace collection, 44 universities all around Germany configured their local caching proxies to use the top level caching proxies of the G-WiN as parents. For NLANR, at time of trace collection 584 individual clients peered to the caches. For both networks, we evaluated the trace files for a period of one week. We present only the results for a single day in the particular week. For all other days recorded in the trace files, similar performance results will be achieved. Trace characteristics for the presented day are shown in Table 1.

Trace	**DFN**	**NLANR**
Date	07/06/2000	6/12/2000
Duration (hours)	24	24
Number of unique documents	2,208,527	1,935,226
Overall size (GB)	42.75	23.07
Number of requests	4,310,753	3,361,744
Requested data (GB)	67.45	64.11
Average requests/sec.	49.89	38.90
Clients	44	584

Table 1. Properties of DFN and NLANR trace

3 Comparative Performance Study

In this section, we study bandwidth consumption and document retrieval latency provided by the different protocols. Due to space limitations, we present only results achieved for the DFN trace.

4 Bandwidth Consumption and Protocol Efficiency

In a first experiment, we study bandwidth consumption in the cache mesh of the local backbone network for the considered cooperative Web caching protocols. As performance measures, we consider saved traffic, i.e., the amount of content delivered from the cache mesh in the local backbone and the inter-cache traffic, i.e., the amount of control messages and data transferred among the Web caches in the local backbone. Figures 2 plots the amount of saved traffic as a function of individual cache size for the DFN trace. Cache sizes range form about 5% to more than 100% of the average volume of data requested form each cache. Not that most misses occur for small cache sizes. Note that for large cache sizes all requested document of the considered traces could be stored in at least one cache of the mesh.

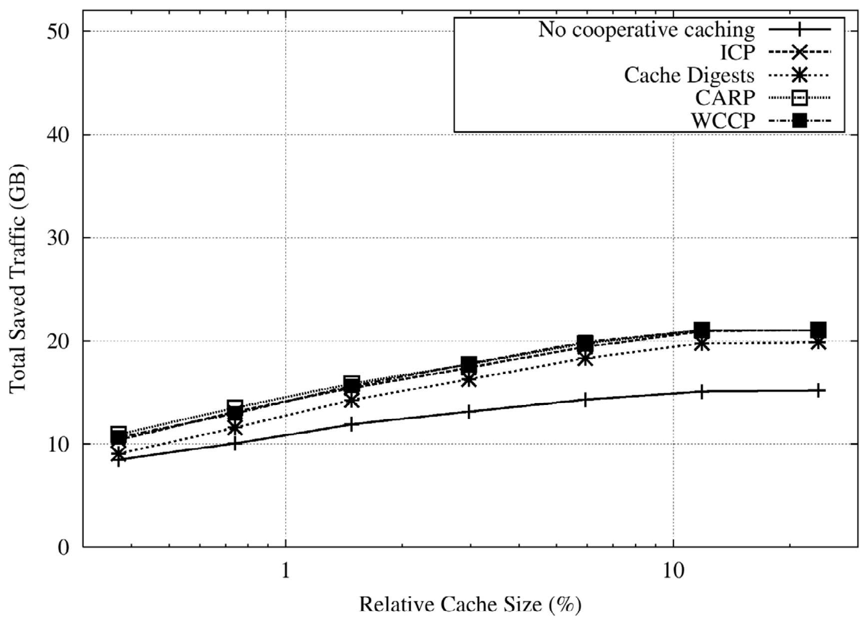

Figure 2. DFN trace: Bandwidth consumption of different protocols; saved traffic

The results show that without cooperative Web caching 22% of the overall requested data volume is saved for DFN trace. Cooperative Web caching increases saved volume up to

29% for the DFN trace. This constitutes an increase of 39% relative to the saved data volume among the caches without cooperative Web caching.

As explained in [9], hash-based schemes achieve best storage efficiency. CARP and WCCP unify several physical caches to one logical cache. That is each document is stored at most one time in the cache mesh and therefore no space is wasted by replicated documents. This leads to best performance in terms of saved volume. An upper bound for achievable performance is the performance of a single caching proxy with combined user population [14].

Consistent for both traces, CARP and WCCP are tightly followed by ICP with respect to bandwidth consumption. Low packet loss probability P_{loss} and low latency for ICP queries and responses enables ICP to discover for almost every request, if a copy of a requested document is stored in the cache mesh before reaching the timeout. Compared to CARP, the saved volume decreases due to local replication of documents. Consistent to [8], we find that ICP successfully implements cooperative Web caching in terms of saved data volume.

The volume saved by Cache Digests is influenced by low validity of the digests. Recall that digests are exchanged on a regular basis, e.g., each cache requests digests from all its peers once an hour. As explained in [10], digests are not synchronized after an exchange. Documents added to a cache after an exchange cannot be localized by examining the digest. Due to this false sharing, a cache must retrieve a document from the original Web server rather than receiving it form a peer.

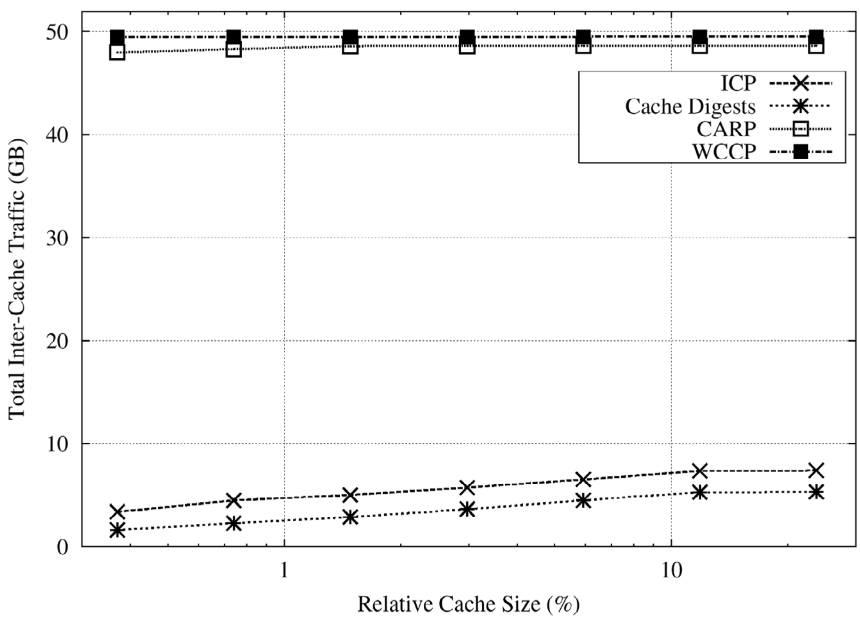

Figure 3. DFN trace: Bandwidth consumption of different protocols; inter-cache traffic

Figure 3 plots volume the inter-cache traffic as a function of cache size for the DFN trace, respectively. Note that without cooperative Web caching no inter-cache communication occurs. Thus, the figure contains only four curves. For CARP and WCCP, the amount for inter-cache communication is about 4/5 of the requested total amount. This can be explained by the fact that the URL space is equally partitioned among the five caches by the hash functions. The Web cache to which a request is issued services only 1/5 of the requested documents. The remaining 4/5 must be transferred from one of the four other caches.

ICP generates UDP traffic for every cache miss in a target cache [13]. It further causes inter-cache traffic when retrieving a document form a peer. UDP traffic decreases with increasing cache size, because hit rates at the target cache increase. In opposite, traffic caused by document retrieval increases because more misses in the target cache can be resolved by hits in the peer caches.

Cache Digests cause the smallest amount of inter-cache communication for two reasons. First, Cache Digests does not introduce communication on every request as CARP or on every miss as ICP. Cache communication is performed by frequently exchange of digests. Second, Cache Digests is not able to discover every copy of a document stored in the cache mesh. Such documents are received from the original Web server, decreasing both inter cache traffic and saved volume.

4.1 Document Retrieval Latency

Figures 4 and 5 provide a comparison of the latencies introduced by the protocols for different bandwidth in the local backbone network. For very small caches, no cooperative Web caching yields best performance in terms of latency in DFN trace. This can be explained considering the probability for a hit at a particular cache inside the cache mesh. If the target cache misses on a document request, the probability for finding the requested document in any other cache inside the mesh is relative small, and so is the expected benefit for trying to receive the document form inside the hierarchy. The benefit is even smaller than overhead introduced by the protocols for cooperative web caching. With growing cache size, cache hit rate grows and so does the expected benefit of cooperative Web caching. For low round trip times, even CARP gets superior to not employing a cooperative Web caching scheme in terms of document retrieval latency. WCCP as a router-based protocol takes the responsibility for request distribution away form the caches. As a consequence, each request is forwarded to exactly one cache. This reduces queuing delays at the caches, which results in low document latency. In terms of latency, WCCP outperforms all protocols but ICP. We conclude from our studies that high Speed backbone networks enable hash-based protocols to be successfully employed for wide area distributed caching.

Comparison between Figures 4 and 5 shows that network round trip time has little impact on document retrieval latency, which is dominated by download delay for access to the original Web server. Faster networks make little impact on efficiency of distributed caching. Latency shrinks linear in RTT.

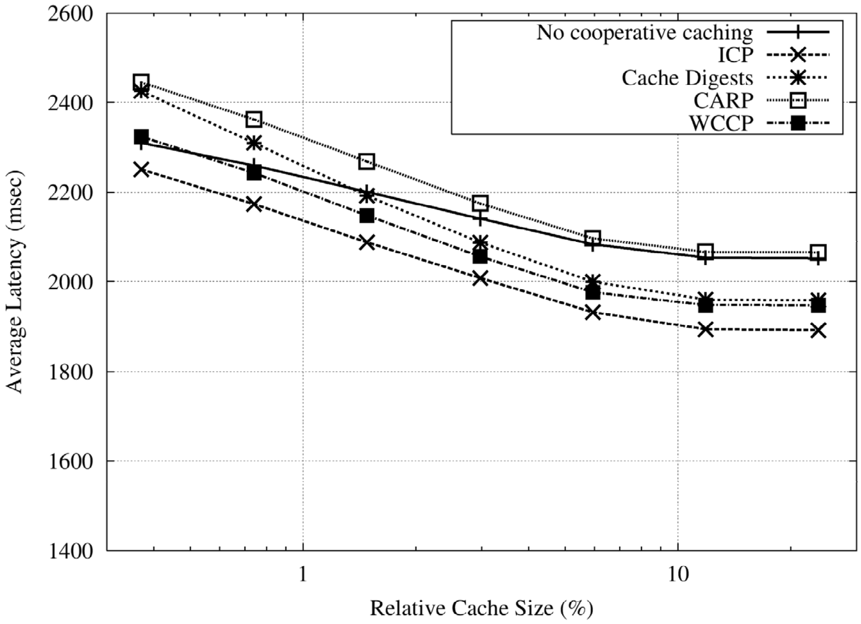

Figure 4. DFN trace: Document latency of different protocols; 155 Mbps bandwidth

5 Conclusion

We presented a comprehensive performance study of the cooperative Web caching protocols ICP, CARP, Cache Digests, and WCCP. The study is based on the same IP backbone network topology and workload characteristics. We investigate cooperative Web caching from viewpoints of the client and the network. As performance measures, we consider user latency and bandwidth consumption. The goal of our study lies in understanding the behavior and limiting factors of cooperative Web caching protocols network for reducing Web document access latencies in the global Internet.

We conclude from our performance studies that Fast backbone networks allow employment of hash routing protocols as CARP or WCCP for distributed caching in wide area networks. Furthermore, we conclude that emerging network technology has little impact on efficiency of distributed caching, because latency is dominated by download delays. Therefore, bigger latency reduction can be achieved rather by replication of origin servers rather than distributed caching.

References

[1] CA*-Net 3, Canada's Advanced Internet Development Organization. http://www.canarie.ca/advnet/canet3.html

[2] M. Cieslak, D. Foster, G. Tiwana, and R. Wilson, Web Cache Coordination Protocol v2.0. IETF Internet draft, 2000. http://www.ietf.org/internet-drafts/draft-wilson-wrec-wccp-v2-00.txt

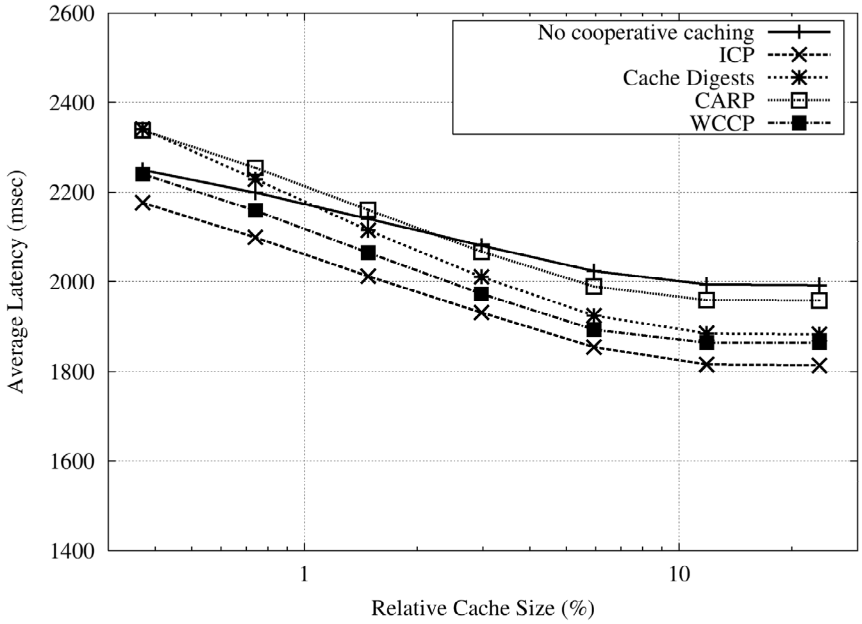

Figure 5. DFN trace: Document latency of different protocols; 2.4 Gbps bandwidth

[3] L. Fan, P. Cao, J. Almeida, and A. Broder, Summary Cache: A Scalable Wide Area Web Cache Sharing Protocol, Proc. ACM Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM), Vancouver Canada, pp. 254-265, 1998.

[4] G-WiN, The German Research Network Association (DFN-Verein), http://www.dfn.de/win/gwin/

[5] Internet-2, University Corporation for Advanced Internet Development. http://www.internet2.edu/

[6] A. Mahanti, C. Williamson, and D. Eager, Workload Characterization of a Web Proxy Caching Hierarchy, IEEE Network Magazine: Special Issue on Web Performance, 14, pp. 16-32, 2000.

[7] NLANR, National Laboratory for Applied Network Research, http://www.nlanr.net/

[8] P. Rodriguez, C. Spanner, and E.W. Biersack, Analysis of Web Caching Architectures: Hierarchical and Distributed Caching, submitted to IEEE/ACM Trans. on Networking 1999, http://www.eurecom.fr/~btroup/FormerPages/rodrigue_htdoc/

[9] K.W. Ross, Hash Routing for Collections of Shared Web Caches, IEEE Network, 11, pp. 37-44, 1997.

[10] A. Rousskov and D. Wessels, Cache Digests, Computer Networks and ISDN Systems, 30, pp. 22-23, 1998.

[11] H. Schwetman, Object-oriented Simulation Modeling with C++/CSIM17, Proc. of the 1995 Winter Simulation Conference, Eds. C. Alexopoulos, K. Kang, W. Lilegdon, D. Goldsman, pp. 529-533, 1995. http://www.mesquite.com

[12] V. Valloppillil and K.W. Ross, Cache Array Routing Protocol v1.0, expired Internet draft, 1998, http://www.ircache.net/Cache/ICP/carp.txt

[13] D. Wessels and K.C. Claffy, ICP and the Squid Web Cache, IEEE Journal on Select. Areas in Comm., 16, pp. 345-357, 1998.

[14] A. Wolman, G. Voelker, N. Sharma, N. Cardwell, A. Karlin, and H. Levy, On the Scale and Performance of Cooperative Web Proxy Caching, Proc. 17th ACM Symp. on Operating Systems Principles (SOSP), Kiawah Island SC, pp. 16-31, 1999.