[Originallly published in the January 2001 issue of *Computing Research News*]

Digital Government Fellow Speaks at SuperComputing 2000

By Geoffrey M. Voelker

At the SuperComputing 2000 conference in November, I presented a talk entitled, "On the Scale and Performance of Cooperative Web Proxy Caching." I gave this talk to inaugurate the Computing Research Association/San Diego Supercomputer Center Digital Government Fellowship Program

The program is supported by the National Science Foundation's Digital Government Program, and is intended to build ties between academic and industrial computing research communities, as well as among information technology workers in federal, state, and local governments. In this article, I describe research to investigate techniques for improving the performance of large-scale, wide-area Web-based information services such as those envisioned in the Digital Government Program.

As described in the report from the 1998 Digital Government Workshop [1], the NSF Digital Government Program is a very broad, multidisciplinary initiative to foster research and collaboration for using information technology to transform public services at all levels of government. Whether digital government results in transformations in public service as convenient as downloading tax forms from the IRS website, as immediate as viewing webcasts of city council meetings or congressional hearings, or as sensational as electronic voting in presidential elections, the systems supporting current and future digital government services will continue to grow in complexity and scale. Given the increasing interdependence among government agencies at the local, state, and federal levels, as well as among the public, private, and nonprofit sectors, the systems providing digital government services will be distributed among many organizations and geographic sites spanning the national internetwork. Furthermore, as government services move online, citizens expect those distributed government services to have the same quality, performance, and convenience of private online services. Given the complexities of the systems providing digital government services, delivering performance for these systems is a key challenge.

One goal of my research is to investigate techniques for improving the performance of large-scale, wide-area Web-based information services such as those envisioned in the Digital Government Program. In particular, the research that I presented in the talk explores the potential for large-scale cooperative Web caching systems to improve Web performance. This research was in collaboration with Alec Wolman, Hank Levy, Anna Karlin and others at the University of Washington.

Internet proxy caching has become a commonplace approach for improving the performance of Web browsers. Typically, the proxy sits in front of an entire company or organization. By caching requests for a group of users, a proxy can quickly return documents previously accessed by other clients. Ultimately, though, the hit rate of the proxy is a function of the size of the population it manages—a size often dictated by political, organizational, or geographic considerations. A compelling question, then, is whether multiple proxies should cooperate with each other in order to increase total client population, improve hit ratios, and reduce document-access latency. Whether such cooperative proxy caching is a useful architecture for improving performance depends on a number of factors. These factors include the sharing patterns of documents across organizations, the ratio of inter-proxy communication time to server fetch time, and the scale at which cooperation is undertaken.

Several cooperative Web caching architectures and protocols have been proposed in the research community for building large-scale distributed caching systems. However, few studies have examined cooperative Web caching from a systemic viewpoint. As a result, we know neither the environments in which cooperative caching is useful (if any) nor its potential performance benefits. Answering such questions has been difficult in the past because studying proxy cooperation requires simultaneous traces from multiple proxies.

In our work, we take a two-pronged approach to exploring the limits and potentials of cooperative proxy caching. As the first approach, we collect and analyze traces from two environments: the University of Washington and the Microsoft Corporation. These traces enable us to investigate the performance benefits of cooperative Web caching systems at small

(100 to 1000s of users) and medium (1000s to 10000s of users) population sizes. As a key component of our university trace, we
identify each client in terms of its membership in one of about 200 university departments or programs. This gives us the equivalent of a simultaneous trace of 200 diverse, independent organizations, permitting us to analyze document sharing among those organizations and to measure the potential benefits of cooperation among organization-based proxies. We examine latency and bandwidth benefits of proxy caching for this data, as well.

We then use the Microsoft trace of employee traffic to the Internet to explore the potential of cooperation between larger organizations. To do this, we analyze traces from Microsoft and the university that we collected over the same time period and processed with the same anonymization function. This permits a direct computation of the degree of document sharing, and hence the benefit of sharing, between two proxies each handling tens of thousands of clients.

As the second approach, we develop an analytic model of Web behavior that extends beyond the limits of our trace results. The model determines the steady-state performance of Web caches as a function of the size of client populations and other workload parameters, such as document rate of change. The model permits us to examine the impact of larger population sizes, to explore the tradeoffs among various cooperative-caching schemes, and to speculate on the performance implications of future trends.

Based on our analyses, we find that the usefulness of cooperative Web proxy caching fundamentally depends on the scale at which it is being applied. From our trace data of users at the University of Washington and the Microsoft Corporation, our results show that cooperative Web proxy caching is an effective architecture for small organizational caches that together comprise user populations in the tens of thousands. Furthermore, given the small scale, short latencies, and high bandwidths, any reasonable cooperative caching scheme will serve.

However, the workload generated by client populations at this scale is not very demanding, and can be easily handled by the use of a single Web proxy cache. As a result, if

the organizations can be pooled behind a single cache, there is no technical motivation for using cooperative caching at this scale; if it is administratively and politically feasible, a single proxy cache can provide the same benefits with fewer resources and less overhead. If, however, organizational, political, or geographic constraints require the use of individual caches for each organization at this scale, then those organizations would significantly improve their performance by having their caches cooperate.

Whether or not they use cooperative caching internally, large organizations such as the University of Washington and Microsoft should use proxy caching for their user populations. A key question, then, is whether these large organizational caches benefit from cooperation. To answer this question, we used our analytic model to explore the performance benefits of cache cooperation among large-scale client populations, and the effect of cooperative cache architecture on performance.

Experiments with our analytic model indicate that cooperation among the organizational caches within a medium to large city will still provide some benefit, although incremental, over cooperative caching at small scales. Assuming that bandwidth within a city is plentiful and latencies are small, the overhead of cooperative caching would be low and potentially worth the secondary benefits that such caching provides. In principle, the organizational caches within a city can use a hash-based scheme to maximize storage efficiency. In practice, however, given the cheap cost of disks, using a hash-based scheme to spread load is more important than storage efficiency. Extrapolating to yet larger scales, such as the state level and even the west coast of the United States, our model results indicate that cooperative caching among cities would provide very limited additional benefit, particularly given the increased latencies among caches.

To summarize, our results do show benefits of cooperative caching among collections of small organizations. However, we show that cooperative caching is unlikely to have significant benefits for larger organizations or populations. That is, with current sharing patterns, there is no motivation for designing highly scalable cooperative-caching schemes; all reasonable schemes will have similar performance in the low-end population range where cooperative caching works.

[1] Sharon S. Dawes et al., "Some Assembly Required: Building a Digital Government for the 21st Century." Report of the 1998 Multidisciplinary Workshop on Digital Government, Center for Technology in Government of the University at Albany/SUNY.

Geoffrey Voelker is an Assistant Professor in the Department of Computer Science at UC San Diego (voelker@cs.ucsd.edu).