Performance and Reliability Issues in Client-Server Environments

Rick Bunt and Derek Eager
Distributed Systems Performance Laboratory
Department of Computational Science
University of Saskatchewan
Saskatoon, Canada S7N 0W0

September 11, 1994

1 Preamble

The client-server environment has supplanted the traditional mainframe environment in many organizations. While this new paradigm presents many exciting computational opportunities, it also gives rise to new problems concerning performance and reliability. As both workstation clients and servers become increasingly well-resourced, and applications become increasingly demanding, a number of system design decisions need to be rethought if the service expectations of the applications are to be met. Our research addresses issues relating to the performance and reliability of client-server environments.

2 Outline of Current Research Activities

The timely availability of data stored in files is essential to the success of any application, and the client-server paradigm presents new challenges in file system management. Some of our current work is considering the problem of disk block caching in a client-server environment. To reduce the delays arising from synchronous I/O operations, caching is introduced in the hope that future file system requests can be resolved in a high-speed cache, thereby avoiding access to a slower medium. Workstation memories are now large enough to accomodate large disk block caches with low miss ratios, and continued technology advances should make it possible to achieve even lower miss ratios in the future. Cache can be located at either the client or the server (or both): at the client to avoid network access to the server, or at the file server to avoid physical disk accesses. The presence of caches at *both* the client and the server, however, results in a two-level hierarchical cache, the performance of which is not sufficiently understood. For example, client caching significantly alters the nature of demand at the file server, with potentially profound performance consequences. Ongoing work concerns the investigation of the impact of various configuration decisions on overall system performance, including the amount of file cache at clients and servers, the cache management policies at each location, and the interaction semantics between the two levels of cache.

Several recent pieces of work have dealt with various aspects of this issue and provide performance comparisons derived from trace-driven simulation experiments. The effect of client caching was first addressed in [Makaroff and Eager 90]. Subsequent work reported in [Willick *et al.* 92] addressed the choice of cache management policy at the server in the light of the removal of temporal locality from the stream of references presented to the server. Although conventional locality-based cache management strategies such as LRU work well at the client it was shown that they perform poorly at the server. Frequency-based approaches such as LFU that are not locality sensitive were found to perform better at the server. A paper presently in preparation [Bunt *et al.* 94] takes a more complete look at the problem and considers the effect of multiple clients as well. The interleaving of request streams from multiple clients exacerbates the problems at the server. Again, locality-based approaches are found to be inappropriate for server cache management.

Typically it has been read operations that benefit most from caching. A recently-completed thesis [Chen 94] looks at write caching using newly emerging non-volatile memory technologies for the cache. When cache is volatile the effects of updates (writes) to cached data must be reflected at the server very soon (e.g. a *write through* approach) to ensure the integrity of the data. Unlike ordinary (volatile) memory, the contents of non-volatile memory are preserved even in the event of power failure. This increases reliability but also makes it possible to delay the write-back operations for cached data. In comparison to write-through based systems, the response time for write operations can be greatly reduced as data need only be written into the non-volatile cache (at memory speed) rather than to the server and to disk (as in write-through based approaches) before control returns to the user program. Furthermore, there are other performance gains possible by delaying the write backs: first, while data is resident in the cache, a significant fraction of it may be overwritten or deleted, thus reducing the amount of write traffic; second, it is possible to amortize the cost of physical disk accesses over multiple write (and read) requests. Our simulations showed that locality-based policies once again performed well for the client write caches, but policies that amortize disk access overhead over multiple I/O operations performed much better at the server.

In general terms, our research deals with the system-level support needed to develop and manage future distributed applications in a client-server environment. The overall objective of this research is to provide an effective and reliable systems platform upon which distributed applications can be developed and run so that their service expectations are met. We are especially interested in emerging multimedia applications which may present substantially different resource demands, demands that are not yet well-understood. Topics of interest to us include workload characterization for multimedia applications and design issues for multimedia file servers. We also have work underway that is looking at issues that arise when clients and servers become temporarily disconnected from each other, either non-voluntarily or voluntarily. What sort of system support is required to maintain some level of operability through periods of disconnection?

As is evident from the work we have done in client-server caching, our research is highly experimental in its approach. Measurement is a key component of any experimental investigation. A central element of this research concerns the problems of performance measurement in file server/workstation clusters, and of understanding and characterizing the workload of these systems. Recent results include a workload model for use in the performance assessment of NFS file servers [Bodnarchuk and Bunt 91] and a model for characterizing the injection and locality patterns of network traffic [Ling *et al.* 91, Gulati *et al.* 93]. Additional measurement tools are needed to support data gathering in a fully distributed setting, supporting both client-server and peer-to-peer models, not only from multiple clients but also from multiple processes executing on the same client. An approach based on the dissemination of data collection agents throughout the distributed environment appears to offer considerable promise.

References

[Bodnarchuk and Bunt 91] Bodnarchuk, R. and Bunt, R.B., "A Synthetic Workload Model for a Distributed System File Server", *Proc. 1991 ACM SIGMETRICS Conference on Measurement and Modeling of Computing Systems*, San Diego, CA, May 1991, 50-59.

[Bunt *et al.* 94] Bunt, R.B., Eager, D.L., Froese, K.W. and Willick, D.L., "The Performance of Disk Caching Strategies in a Distributed File System", in preparation.

[Chen 94] Chen, K., *Write Caching in Distributed File Systems*, M.Sc. Thesis, Department of Computational Science, University of Saskatchewan, July 1994.

[Gulati *et al.* 93] Gulati, N., Williamson, C.L. and Bunt, R.B., "Local Area Network Traffic Locality: Characterization and Application", *Proc. First International*

Conference on LAN Interconnection, Raleigh, NC, October 1993, 233-250 [Proceedings published as *Local Area Network Interconnection* (eds. R.O. Onvural and R. Nilsson), Plenum Press, 1993].

[Ling *et al.* 91] Ling, S.H, Bunt, R.B. and Eager, D.L., "Characterising Client-Server Workload on an Ethernet", *Proc. CMG '91 Conference*, Nashville, TN, December 1991, 59-69.

[Makaroff and Eager 90] Makaroff, D.J. and Eager, D.L., "Disk Cache Performance for Distributed Systems", *Proc. 10th International Conference on Distributed Computing Systems*, Paris, France, May 1990, 212-219.

[Willick *et al.* 92] Willick, D.L., Eager, D.L. and Bunt, R.B., "Disk Cache Replacement Policies for Network Fileservers", *Proc. 13th International Conference on Distributed Computing Systems*, Pittsburgh, PA, May 1993, 2-11.