

CITI Technical Report 91-3

Multi-level Caching in Distributed File Systems

D. Muntz

dmuntz@citi.umich.edu

P. Honeyman

honey@citi.umich.edu

ABSTRACT

We are investigating the potential for intermediate file servers to address scaling problems in increasingly large distributed file systems. To this end, we have run trace-driven simulations based on data from DEC-SRC and our own data collection to determine the potential of caching-only intermediate servers.

The degree of sharing among clients is central to the effectiveness of an intermediate server. This turns out to be quite low in the traces available to us. All told, fewer than 10% of block accesses are to files shared by more than one file system client.

Trace-driven simulation shows that even with an infinite cache at the intermediate, cache hit rates are disappointingly low. For client caches as small as 20 MB, we observe hit rates less than 19%. As client cache sizes increase, the hit rate at the intermediate approaches the degree of sharing among all clients. On the other hand, the intermediate does appear to be effective in reducing the peak load presented to upstream file servers.

August 16, 1991

Multi-level Caching in Distributed File Systems

D. Muntz

dmuntz@citi.umich.edu

P. Honeyman

honey@citi.umich.edu

Introduction

As distributed file systems grow, so does the need to increase scalability. At the Institutional File System Project, we are investigating tools and techniques for offering file service to a huge client base, perhaps as many as 30,000 end systems.

One problem encountered along the way is the need to service clients supporting a variety of file system protocols, *e.g.*, AFS [1], NFS [2], and AFP [3]. Our principal file servers all run AFS, so the first of these is not a problem. For other file system protocols, we have built intermediate servers that act as AFS clients of the principal file servers and as NFS or AFP servers for clients using these other protocols.

We also considered the case where the intermediate server uses AFS for both the client and server protocols. This architecture extends to one in which there are multiple levels of intermediate AFS (or *iAFS*) servers, each caching files it fetches from the upstream servers, and serving files out of its cache to downstream clients.

One reason for considering multi-level cache hierarchies is that they have shown great success in improving CPU performance when used in processor memories [4]. In the context of file systems, intermediate caches have the potential to increase the capacity of upstream file servers by reducing the number of requests that they must process and the amount of state information that they must maintain. The goal of this study is to assess that potential. Our principal tool is a trace-driven simulator that analyzes file system trace data taken from ‘‘real-world’’ networks.

Trace-driven simulation

To explore the potential of multi-level caching in distributed file systems, we ran trace-driven simulations to predict the hit rates that we might see at an iAFS server. The traces fed to the simulator were derived both from data collected in a network of Firefly workstations [5] at the Digital Equipment Corporation’s Systems Research Center, and from file server trace data collected here.

Firefly trace data

The Firefly data was collected over a four day period in February, 1990 from 115 Firefly workstations supporting the Topaz environment, which includes a (proprietary) distributed file system protocol. During the trace period, each client produced a log record for every system call related to file system operations. Each record contained the following information:

- the name of the system call
- the process id of the invoking process
- the arguments to the call
- the time at which the call was entered
- the time at which the call was exited
- the success or failure status of the call

We preprocessed the data to convert file descriptors into path names, to eliminate irrelevant log records, and to normalize the name space.

The cache simulator expects pathnames in the trace stream to account for hits and misses. However, some system calls, *e.g.*, `read`, use a file descriptor instead of a pathname. To convert fd’s to pathnames, we implemented a process simulator that builds a table for each process which associates the pathname used in, say, `open` calls with the file descriptor returned. The table is

copied across fork and exec calls. Relative pathnames, such as those starting with “.” and “..”, as well as those containing “.” and “..”, were also converted to the appropriate pathnames at this stage.

In this study, accesses to the local file system were not of interest and were eliminated in preprocessing. In addition, system calls that failed were elided. Failures can arise, *e.g.*, when attempting to create a file in a write-protected directory.

The name space was normalized by converting names of the form `host:path` to a flat name space of unique integers. In all, 68,413 different pathnames are referenced in 2,807,003 trace records.

IFS trace data


The IFS trace data was collected from four AFS servers running on IBM RT computers during a 4.8 day period in early November, 1990. The data records all file server requests from 49 clients. 31,538 different files are referenced in 92,571 trace records.

Data collected at the IFS Project was tailored more directly to our needs and required less preprocessing. AFS clients exchange file data with the server via `FETCHDATA` and `STOREDATA` requests, whose functions follow directly from their names. Each `FETCHDATA` and `STOREDATA` request contains a timestamp, the client's network address, the file's `FID` (the unique identifier for a file) and the offset and length of the data being requested.

In the IFS traces the return status of the requests is not recorded; we have observed that almost all fetch and store requests succeed, so we don't believe that this limitation invalidates the results reported by the simulator.

The simulator

We performed experiments simulating distributed environments with a two-level cache design, using file system activity traces provided by DEC-SRC as well as traces collected locally. In the simulated environments of the experiments, the client machines are connected to an intermediate server which is in turn connected to a “main” server.



For the experiments discussed in this paper, the intermediate server has a potentially infinite cache. This is obviously impractical. Because an iAFS server with a finite cache would be forced to flush its contents on occasion, the hit rates reported here are larger than can be achieved in reality.

The operation of the simulator is straightforward. When a system call requesting a file from the server machine appears in the trace data, the simulator checks the local cache on the requesting machine to see if the request can be satisfied there. If the requested block is found in the local cache, a “hit” is logged for that client and the next trace record is processed. Otherwise, a “miss” is recorded for the client, and the cache on the iAFS server is checked for the requested block.

If the block is found in the iAFS server's cache, a hit is recorded for the iAFS server, and the block is placed in the client's cache. Otherwise, a miss is recorded for the iAFS server, the block is installed in both the iAFS server's cache and the client's cache, and the next trace record is processed.

In this way, the input of trace records is processed until exhausted. All read and write requests are guaranteed to succeed at the server, and the cache replacement policy is LRU. When a block is written, the simulator invalidates that block in the requesting client's cache.

Hit rate simulations


In the first set of experiments, we examine the hit rates that can be expected for the iAFS server cache. We first simulate an iAFS server with an unbounded cache using the trace data from the Firefly clients. We then restrict our attention to 20 Firefly clients that appear to exhibit a high degree of data sharing — these 20 clients are

responsible for over half of the iAFS server cache hits. We then use the trace data collected from the 49 IFS clients. Again we simulate an iAFS server with an unbounded cache.

We used 64K as our cache block size, because this is the size used by AFS. Simulations were also run using block sizes of 4K, 8K, 16K, and 32K; those results are not substantially different from the ones presented here.

Firefly clients

The first experiment with the Firefly data simulates an environment in which all 115 machines are clients to an iAFS server. The iAFS server is given an “infinite” cache size, so that if a given block is ever sought twice, each request after the first causes a hit at the iAFS server. In practice, the iAFS server cache would have to be 7,880M to achieve this hit rate. The size of the client caches is varied in each simulation to generate a graph of client cache size vs. iAFS server cache hit rate.



Client cache size vs. intermediate hit rate


Client caching is desirable: when client caching is disabled, the iAFS has a 70% hit rate, which falls off rapidly as clients are able to resolve requests from a local cache. With just 1M of client cache, the iAFS hit rate falls below 30%. In our (typical AFS) environment, clients have 20M caches; the simulator predicts that an iAFS server would show an 11% hit rate. Client cache sizes can be expected to grow as disk density continues to increase. The simulation predicts a corresponding decrease in iAFS server hit rates: an 80M client cache produces a 7% iAFS hit rate.

As the client cache size approaches infinity, the hit rate at the intermediate asymptotically approaches the degree of sharing. This asymptote, the degree of sharing among all clients, is represented in the graphs by a dashed line. The degree of sharing among the Firefly clients is

6.1%.

Partial Firefly clients

Among the Firefly clients, there is a subset whose file reference patterns are more tightly woven: 20 clients are responsible for over half of the overlap in file references among all 115 Fireflies. We simulated an iAFS server for these 20 clients, as in the previous section. The degree of sharing among the 20 is 5.8%.




Client cache size vs. intermediate hit rate

Again, we see that even meager client cache sizes result in low intermediate hit rates. In this experiment when the client caches are 20M, the intermediate hit rate is about 18%. When client caches are 80M, the intermediate hit rate drops to about 7.6%

IFS clients

In the third experiment, data collected on IFS project file servers was used to drive the simulations. The four servers on which data was collected contain all home directories, system binaries, and project-related data and programs for the several dozen IFS project staff.

This data contains `FETCHDATA` and `STOREDATA` requests from 49 AFS clients. Again, the simulation involves all 49 clients connected to one intermediate with infinite cache. Clients in the IFS project have 20M caches on their local disks, and `FETCHDATA` requests satisfied by the local cache do not appear in traces recorded at the server. Because the trace data was collected on the server, simulation is possible only for client caches of at least 20M. The degree of sharing among the 49 IFS project clients is 8.5%.




Client cache size vs. intermediate hit rate

For typical AFS client cache sizes, 20–80M, the simulator predicts the iAFS hit rate to be between 15% and 18%.


Hot and cold cache experiments

In this section, we focus on the IFS trace data with 20M client caches. In the simulations described so far, the iAFS server cache is initially empty, or “cold.” Consequently, in the early hours of the simulation there are very few cache hits. Examining the iAFS server hit rate over time, the aggregate hit rate increases through the simulated 116 hour duration, and is apparently still rising at the end of the simulated period.



Time vs. iAFS server cache hit rate


As the iAFS server cache “warms up,” the hit rate becomes more respectable. Clearly there must be times when the hit rate is higher than the final 18%. In the next experiment, we collect hourly hit rate statistics and plot them as the “instantaneous” hit rate at the iAFS server cache. We show this in the next graph, along with the hourly request rate presented to the iAFS server.



Upper graph: instantaneous hit rate
Lower graph: request rate

The upper graph shows a highly variable instantaneous hit rate with several distinct peaks. The lower graph shows a similar pattern in request rates. Note that after a “warm-up” interval, the peaks in the request rate coincide with peaks in the hit rate.

To gauge the effect of a “hot” cache in the iAFS server, we re-ran the simulation of the previous section on a hot-cache iAFS server. We treat the first half of the simulation period, 58 hours, as the warm-up interval and gather statistics for various client cache sizes in the last half of the period. The following graph shows the results of this experiment, superimposed with the graph from the preceding section, where warm-up is not taken into consideration.



Client cache size vs. intermediate hit rate


The upper graph shows an improvement in the iAFS server cache hit rate when it is pre-heated in the first 58 hours of the trace interval. Even here, though, the iAFS hit rate is below 30% for standard 20M client caches.

Another set of experiments involves clearing the client caches periodically while maintaining accounting throughout the simulation. This models an environment in which a machine is used sequentially by different people. The results

are more pessimistic, however, as there would likely be some overlap among the users' data requests *e.g.*, `/bin/csh`. Results from these tests are similar to the other warm cache experiments: the simulation predicts an improved hit rate at the iAFS server, but the improvement is not dramatic.


Effect on upstream server load

In the previous section, we saw that high request rates appear to correlate with high instantaneous hit rates on the iAFS server. This suggests that the iAFS server may be effective in moderating the peak traffic presented to the upstream server(s). To test this hypothesis, we ran a set of simulation experiments with the IFS data to measure the request rate seen by upstream servers when an iAFS server is present and when it is absent.




The graph shows the request rate, in requests per hour, presented to upstream server(s) when the iAFS server is present (solid line) and when it is absent (dotted line). After an interval during which the iAFS server cache warms up, the effect is striking: the peak load is reduced from over 2,500 requests per hour to fewer than 1,400 requests per hour.

The Firefly data also shows a correlation between request rates and iAFS server cache hit rates. Our simulations predict that with those file system traces as well, the iAFS is effective in clipping the peak load presented to the upstream server(s).



The spike occurring late Thursday night is caused by a system task running on a single machine. This task maintains a database of cross-references between pieces of software at DEC-SRC, *e.g.*, which components use which other components. This task accesses a significant portion of the file system. Eliminating this process' activity from the traces, which accounts for about 3% of the Firefly trace data, makes it easier to see the effect of the iAFS on off-loading the upstream server.



Discussion

The simulated hit rates on the iAFS server do not lend much encouragement for its role in enhancing client performance. Our simulations indicate that most of the requests presented to an iAFS server must be forwarded to an upstream server to be satisfied; from a client's perspective, the iAFS can be viewed as a "delay server."

Simulations using the Firefly data show that an iAFS server cache suffers hit rates below 19% when client caches are 20M or more. Simulations using the IFS trace data also predict iAFS server cache hit rates below 18%. This is largely due to a low degree of sharing among clients, less than

9% in both sets of trace data.

We also simulated several “warm cache” scenarios, in which hit and miss accounting is delayed during a warm-up period. These warm cache simulations predict some, but not much, improvement for the iAFS hit rate.

Our simulations indicate that an iAFS server does help server performance, by clipping the peak request load presented by file system clients. We plan further experiments to investigate this and other ways to exploit multi-level caching in distributed file systems.

Acknowledgements

The Firefly traces were gathered by Andy Hisgen, who kindly made them available to us. Susan Owicki, B. Kumar, Jim Gettys, and Deborah Hwang contributed to the file system tracing facility.

We thank Bill Tetzlaff of IBM Research for suggesting some interesting experiments.

This work was partially supported by IBM.

References

1. J.H. Howard, “An Overview of the Andrew File System,” pp. 23–26 in *Winter 1988 USENIX Conference Proceedings*, Dallas (February, 1988).
2. R. Sandberg, D. Goldberg, S. Kleiman, D. Walsh, and B. Lyon, “Design and Implementation of the Sun Network Filesystem,” pp. 119–130 in *Summer 1985 USENIX Conference Proceedings*, Portland (June, 1985).
3. G.S. Sidhu, R.F. Andrews, and A.B. Oppenheimer, *Inside AppleTalk*, Addison-Wesley, Reading (1989).
4. J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*, Morgan Kaufmann Publishers, Inc., Palo Alto (1990).
5. Charles P. Thacker, Lawrence C. Stewart, and Edwin H. Satterthwaite, Jr., “Firefly: A Multiprocessor Workstation,” *IEEE Transactions on Computers* **37**(8), pp. 909–920 (August, 1988).