

BASE: Using Abstraction to Improve Fault Tolerance

MIGUEL CASTRO

Microsoft Research

and

RODRIGO RODRIGUES and BARBARA LISKOV

MIT Laboratory for Computer Science

Software errors are a major cause of outages and they are increasingly exploited in malicious attacks. Byzantine fault tolerance allows replicated systems to mask some software errors but it is expensive to deploy. This paper describes a replication technique, BASE, which uses abstraction to reduce the cost of Byzantine fault tolerance and to improve its ability to mask software errors. BASE reduces cost because it enables reuse of off-the-shelf service implementations. It improves availability because each replica can be repaired periodically using an abstract view of the state stored by correct replicas, and because each replica can run distinct or nondeterministic service implementations, which reduces the probability of common mode failures. We built an NFS service where each replica can run a different off-the-shelf file system implementation, and an object-oriented database where the replicas ran the same, nondeterministic implementation. These examples suggest that our technique can be used in practice—in both cases, the implementation required only a modest amount of new code, and our performance results indicate that the replicated services perform comparably to the implementations that they reuse.

Categories and Subject Descriptors: C.2.0 [**Computer-Communication Networks**]: General—*Security and protection*; C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Client/server*; D.4.3 [**Operating Systems**]: File Systems Management; D.4.5 [**Operating Systems**]: Reliability—*Fault tolerance*; D.4.6 [**Operating Systems**]: Security and Protection—*Access controls; authentication; cryptographic controls*; D.4.8 [**Operating Systems**]: Performance—*Measurements*; H.2.0 [**Database Management**]: General—*Security, integrity, and protection*; H.2.4 [**Database Management**]: Systems—*Object-oriented databases*

General Terms: Security, Reliability, Algorithms, Performance, Measurement

Additional Key Words and Phrases: Byzantine fault tolerance, state machine replication, proactive recovery, asynchronous systems, N-version programming

This research was partially supported by DARPA under contract F30602-98-1-0237 monitored by the Air Force Research Laboratory. Rodrigo Rodrigues was partially funded by a Praxis XXI fellowship.

Authors' addresses: M. Castro, Microsoft Research, 7 J. J. Thomson Avenue, Cambridge CB3 0FB, UK; email: mcastro@microsoft.com; R. Rodrigues and B. Liskov, MIT Laboratory for Computer Science, 545 Technology Sq., Cambridge, MA 02139; email: {rodrigo,liskov}@lcs.mit.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2003 ACM 0734-2071/03/0800-0236 \$5.00

1. INTRODUCTION

There is a growing demand for highly-available systems that provide correct service without interruptions. These systems must tolerate software errors because they are a major cause of outages [Gray and Siewiorek 1991]. Furthermore, there is an increasing number of malicious attacks that exploit software errors to gain control or deny access to systems that provide important services.

This paper proposes a replication technique, BASE, that combines Byzantine fault tolerance [Pease et al. 1980] with work on data abstraction [Liskov and Guttag 2000]. Byzantine fault tolerance allows a replicated service to tolerate arbitrary behavior from faulty replicas—behavior caused by a software bug or an attack. Abstraction hides implementation details to enable the reuse of off-the-shelf implementations of important services (e.g., file systems, databases, or HTTP daemons) and to improve the ability to mask software errors.

We extended the BFT library [Castro and Liskov 1999, 2000, 2002] to implement BASE. (BASE is an acronym for BFT with Abstract Specification Encapsulation.) The original BFT library provides Byzantine fault tolerance with good performance and strong correctness guarantees if no more than one-third of the replicas fail within a small window of vulnerability. However, it requires all replicas to run the same service implementation and to update their state in a deterministic way. Therefore, it cannot tolerate deterministic software errors that cause all replicas to fail concurrently and it complicates reuse of existing service implementations because it requires extensive modifications to ensure identical values for the state of each replica.

The BASE library and methodology described in this paper correct these problems—they enable replicas to run different or nondeterministic implementations. The methodology is based on the concepts of *abstract specification* and *abstraction function* from work on data abstraction [Liskov and Guttag 2000]. We start by defining a common *abstract specification* for the service, which specifies an *abstract state* and describes how each operation manipulates the state. Then we implement a *conformance wrapper* for each distinct implementation to make it behave according to the common specification. The last step is to implement an *abstraction function* (and one of its inverses) to map from the concrete state of each implementation to the common abstract state (and vice versa).

The methodology offers several important advantages:

- Reuse of existing code.* BASE implements a form of state machine replication [Lampert 1978; Schneider 1990] that allows replication of services that perform arbitrary computations, but requires determinism: all replicas must produce the same sequence of results when they process the same sequence of operations. Most off-the-shelf implementations of services fail to satisfy this condition. For example, many implementations produce timestamps by reading local clocks, which can cause the states of replicas to diverge. The conformance wrapper and the abstract state conversions enable the reuse of existing implementations. Furthermore, these implementations can be nondeterministic, which reduces the probability of common mode failures.

- Software rejuvenation through proactive recovery.* It has been observed [Huang et al. 1995] that there is a correlation between the length of time software runs and the probability that it fails. BASE combines proactive recovery [Castro and Liskov 2000] with abstraction to counter this problem. Replicas are recovered periodically even if there is no reason to suspect they are faulty. Recoveries are staggered so that the service remains available during rejuvenation to enable frequent recoveries. When a replica is recovered, it is rebooted and restarted from a clean state. Then it is brought up to date using a correct copy of the abstract state that is obtained from the group of replicas. Abstraction may improve availability by hiding corrupt concrete states, and it enables proactive recovery when replicas do not run the same code or run code that is nondeterministic.
- Opportunistic N-version programming.* Replication is not useful when there is a strong positive correlation between the failure probabilities of the different replicas, for example, deterministic software bugs cause all replicas to fail at the same time when they run the same code. N-version programming [Chen and Avizienis 1978] exploits design diversity to reduce the probability of correlated failures, but it has several problems [Gray and Siewiorek 1991]: it increases development and maintenance costs by a factor of N or more, adds unacceptable time delays to the implementation, and does not provide a mechanism to repair faulty replicas.

BASE enables an opportunistic form of N-version programming by allowing us to take advantage of distinct, off-the-shelf implementations of common services. This approach overcomes the defects mentioned above: it eliminates the high development and maintenance costs of N-version programming, and also the long time-to-market. Additionally, we can repair faulty replicas by transferring an encoding of the common abstract state from correct replicas.

Opportunistic N-version programming may be a viable option for many common services, for example, relational databases, HTTP daemons, file systems, and operating systems. In all these cases, competition has led to four or more distinct implementations that were developed and are maintained separately but have similar (although not identical) functionality. Since each off-the-shelf implementation is sold to a large number of customers, the vendors can amortize the cost of producing a high quality implementation. Furthermore, the existence of standard protocols that provide identical interfaces to different implementations, e.g., ODBC [Geiger 1995] and NFS [RFC-1094 1989], simplifies our technique and keeps the cost of writing the conformance wrappers and state conversion functions low. We can also leverage the effort towards standardizing data representations using XML [W3C 2000].

This paper explains the methodology by giving two examples, a replicated file service where replicas run different operating systems and file systems, and a replicated object-oriented database, where the replicas run the same implementation but the implementation is nondeterministic. The paper also provides an evaluation of the methodology based on these examples; we evaluate the complexity of the conformance wrapper and state conversion functions and the overhead they introduce.

The remainder of the paper is organized as follows. Section 2 describes our methodology and the BASE library. Section 3 explains how we applied the methodology to build the replicated file system and object-oriented database. We evaluate our technique in Section 4. Section 5 discusses related work and Section 6 presents our conclusions.

2. THE BASE TECHNIQUE

This section provides an overview of our replication technique. It starts by describing the methodology that we use to build a replicated system from existing service implementations. Then it describes the replication algorithm that we use and it ends with a description of the BASE library.

2.1 Methodology

The goal is to build a replicated system by reusing a set of off-the-shelf implementations, I_1, \dots, I_n , of some service. Ideally, we would like n to equal the number of replicas so that each replica can run a different implementation to reduce the probability of simultaneous failures. But the technique is useful even with a single implementation.

Although off-the-shelf implementations of the same service offer roughly the same functionality, they behave differently: they implement different specifications, S_1, \dots, S_n , using different representations of the service state. Even the behavior of different replicas that run the same implementation may be different when the specification they implement is not strong enough to ensure deterministic behavior. For example, the NFS specification [RFC-1094 1989] allows implementations to choose the value of file handles arbitrarily.

BASE, like any form of state machine replication, requires determinism: replicas must produce the same sequence of results when they execute the same sequence of operations. We achieve determinism by defining a *common abstract specification*, S , for the service that is strong enough to ensure deterministic behavior. This specification defines the abstract state, an initial state value, and the behavior of each service operation.

The specification is defined without knowledge of the internals of each implementation. It is sufficient to treat them as black boxes, which is important to enable the use of existing implementations. Additionally, the abstract state captures only what is visible to the client rather than mimicking what is common in the concrete states of the different implementations. This simplifies the abstract state and improves the effectiveness of our software rejuvenation technique.

The next step, is to implement *conformance wrappers*, C_1, \dots, C_n , for each of I_1, \dots, I_n . The conformance wrappers implement the common specification S . The implementation of each wrapper C_i is a veneer that invokes the operations offered by I_i to implement the operations in S ; in implementing these operations this veneer makes use of a *conformance representation* that stores whatever additional information is needed to allow the translation from the concrete behavior of the implementation to the abstract behavior. The conformance wrapper also implements some additional methods

that allow a replica to be shut down and then restarted without loss of information.

The final step is to implement the *abstraction function* and one of its inverses. These functions allow state transfer among the replicas. State transfer is used to repair faulty replicas, and also to bring slow replicas up-to-date when messages they are missing have been garbage collected. For state transfer to work, replicas must agree on the value of the state of the service after executing a sequence of operations; they will not agree on the value of the concrete state but our methodology ensures that they will agree on the value of the abstract state. The abstraction function is used to convert the concrete state stored by a replica into the abstract state, which is transferred to another replica. The receiving replica uses the inverse function to convert the abstract state into its own concrete state representation.

To enable efficient state transfer between replicas, the abstract state is defined as an array of objects. The array has a fixed maximum size, but the objects it contains can vary in size. We explain how this representation enables efficient state transfer in Section 2.3.

There is an important trend that makes it easier to apply the methodology. Market forces pressure vendors to offer interfaces that are compliant with standard specifications for interoperability, for example, ODBC [Geiger 1995]. Usually, a standard specification S' cannot be used as the common specification S because it is too weak to ensure deterministic behavior. But it can be used as a basis for S and, because S and S' are similar, it is relatively easy to implement conformance wrappers and state conversion functions, and these implementations can be reused across implementations. This is illustrated by the replicated file system example in Section 3. In this example, we take advantage of the NFS standard by using the same conformance wrapper and state conversion functions to wrap different implementations.

2.2 The BFT Algorithm

BFT [Castro and Liskov 1999, 2000, 2002] is an algorithm for state machine replication [Lamport 1978; Schneider 1990] that offers both liveness and safety provided at most $\lfloor \frac{n-1}{3} \rfloor$ out of a total of n replicas are faulty. This means that clients eventually receive replies to their requests and those replies are correct according to linearizability [Herlihy and Wing 1987; Castro 2000].

BFT is safe in asynchronous systems like the Internet: it does not rely on any synchrony assumption to provide safety. In particular, it never returns bad replies even in the presence of denial-of-service attacks. Additionally, it guarantees liveness provided message delays are bounded eventually. The service may be unable to return replies when a denial of service attack is active but clients are guaranteed to receive replies when the attack ends.

Since BFT is a state-machine replication algorithm, it has the ability to replicate services with complex operations. This is an important defense against Byzantine-faulty clients: operations can be designed to preserve invariants on the service state, to offer narrow interfaces, and to perform access control. BFT provides safety regardless of the number of faulty clients and the safety

property ensures that faulty clients are unable to break these invariants or bypass access controls.

There is also a *proactive recovery* mechanism for BFT that recovers replicas periodically even if there is no reason to suspect that they are faulty. This allows the replicated system to tolerate any number of faults over the lifetime of the system provided fewer than one-third of the replicas become faulty within a window of vulnerability.

The basic idea in BFT is simple. Clients send requests to execute operations and all nonfaulty replicas execute the same operations in the same order. Since replicas are deterministic and start in the same state, all nonfaulty replicas send replies with identical results for each operation. The client chooses the result that appears in at least $f + 1$ replies. The hard problem is ensuring non-faulty replicas execute the same requests in the same order. BFT uses a combination of primary backup and quorum replication techniques to order requests. Replicas move through a succession of numbered configurations called *views*. In a view, one replica is the *primary* and the others are *backups*. The primary picks the execution order by proposing a sequence number for each request. Since the primary may be faulty, the backups check the sequence numbers and trigger *view changes* to select a new primary when it appears that the current one has failed.

BFT incorporates a number of important optimizations that allow the algorithm to perform well so that it can be used in practice. The most important optimization is the use of symmetric cryptography to authenticate messages. Public-key cryptography is used only to exchange the symmetric keys. Other optimizations reduce the communication overhead: the algorithm uses only one message round trip to execute read-only operations and two to execute read-write operations, and it uses batching under load to amortize the protocol overhead for read-write operations over many requests. The algorithm also uses IP multicast and other optimizations to reduce protocol overhead as the operation argument and result sizes increase. Additionally, it provides efficient techniques to garbage collect protocol information, and to transfer state to bring replicas up-to-date.

BFT has been implemented as a generic program library with a simple interface. The BFT library can be used to provide Byzantine-fault-tolerant versions of different services. It has been used to implement Byzantine-fault-tolerant distributed file system, BFS, which supports the NFS protocol. A detailed performance evaluation of the BFT library and BFS appears in Castro [2000] and Castro and Liskov [2002]. This evaluation includes results of micro-benchmarks to measure performance in the normal case, during view changes, and during state transfers. It also includes experiments to evaluate the impact of each optimization used by BFT. Furthermore, these experimental results are backed by an analytical performance model [Castro 2000].

2.3 The BASE Library

The BASE library extends BFT with the features necessary to support the methodology. The BFT library requires all replicas to run the same service

```

Client call:
int invoke(Byz_req *req, Byz_rep *rep, bool read_only);

Execution upcall:
int execute(Byz_req *req, Byz_rep *rep, int client, Byz_buffer *non-det);

State conversion upcalls:
int get_obj(int i, char** obj);

void put_objs(int nobjs, char **objs, int *indices, int *sizes);

Checkpointing:
void modify(int nobjs, int* objs);

Non-determinism upcalls:
int propose_value(Seqno seqno, Byz_buffer *req, Byz_buffer *non-det);

int check_value(Seqno seqno, Byz_buffer *req, Byz_buffer *non-det);

Recovery upcalls:
void shutdown_proc(FILE *out);

void restart_proc(FILE *in);

```

Fig. 1. BASE interface and upcalls.

implementation and to update their state in a deterministic way, which complicates reuse of existing service implementations because it requires extensive modifications to ensure identical values for the state of each replica. The BASE library corrects these problems. Figure 1 presents its interface.

The `invoke` procedure is called by the client to invoke an operation on the replicated service. This procedure carries out the client side of the replication protocol and returns the result when enough replicas have responded.

When the library needs to execute an operation at a replica, it makes an upcall to an `execute` procedure that is implemented by the conformance wrapper for the service implementation run by the replica.

State transfer. To perform state transfer in the presence of Byzantine faults, it is necessary to be able to prove that the state being transferred is correct. Otherwise, faulty replicas could corrupt the state of out-of-date but correct replicas. (A detailed discussion of this point can be found in Castro and Liskov [2000].) Consequently, replicas cannot discard a copy of the state produced after executing a request until they know that the state produced by executing later requests can be proven correct. Replicas could keep a copy of the state after executing each request but this would be too expensive. Instead replicas keep just the current version of the concrete state plus copies of the abstract state produced every k -th request (e.g., $k = 128$). These copies are called checkpoints. Replicas inform each other when they produce a checkpoint and the library only transfers checkpoints between replicas.

Creating checkpoints by making full copies of the abstract state would be too expensive. Instead, the library uses copy-on-write such that checkpoints only contain the differences relative to the current abstract state. Similarly,

transferring a complete checkpoint to bring a recovering or out-of-date replica up to date would be too expensive. The library employs a hierarchical state partition scheme to transfer state efficiently. When a replica is fetching state, it recurses down a hierarchy of meta-data to determine which partitions are out-of-date. When it reaches the leaves of the hierarchy (which are the abstract objects), it fetches only the objects that are corrupt or out-of-date. This is described in detail in Castro and Liskov [2000].

As mentioned earlier, to implement checkpointing and state transfer efficiently, we require that the abstract state be encoded as an array of objects, where the objects can have variable size. This representation allows state transfer to be done on just those objects that are out-of-date or corrupt.

The current implementation of the BASE library requires the array to have a fixed size. This limits flexibility in the definition of encodings for the abstract state but it is not an intrinsic problem. The maximum number of entries in the array can be set to an extremely large value without allocating extra space for the portion of the array that is not used, and without degrading the performance of state transfer and checking.

To implement state transfer, each replica must provide the library with two upcalls, which implement the *abstraction function* and one of its inverses. These upcalls do not convert the entire state each time they are called because this would be too expensive. Instead, they perform conversions at the granularity of an object in the abstract state array. The abstraction function is implemented by `get_obj`. It receives an object index i , allocates a buffer, obtains the value of the abstract object with index i , and places that value in the buffer. It returns the size for that object and a pointer to the buffer.

The inverse abstraction function receives a new abstract state value and updates the concrete state to match this argument. This function should also work incrementally to achieve good performance. But it cannot process just one abstract object per invocation because there may be invariants on the abstract state that create dependencies between objects. For example, suppose that an object in the abstract state of a file system can be either a file or a directory. If a slow replica misses the operations that create a directory, d , and a file, f , in d , it has to fetch the abstract objects corresponding to d and f from the others. Then, it invokes the inverse abstraction function to bring its concrete state up-to-date. If f is the argument to the first invocation and d is the argument to the second, it is impossible for the first invocation to update the concrete state because it has no information on where to create the file. The reverse order does not work either because the first invocation creates a dangling reference in d .

To solve this problem, `put_objs` receives a vector of objects with the corresponding sizes and indices in the abstract state array. The library guarantees that this upcall is invoked with an argument that brings the abstract state of the replica to a consistent value (i.e., the value of a valid checkpoint).

Each time the `execute` upcall is about to modify an object in the abstract state it is required to invoke a `modify` procedure, which is supplied by the library, passing the object index as argument. This is used to implement copy-on-write to create checkpoints incrementally. When `modify` is invoked, the library checks if it has saved a copy of the object since the last checkpoint was taken.

If it has not, it calls `get_obj` with the appropriate index and saves the copy of the object until the corresponding checkpoint can be discarded. It may be difficult to identify which parts of the service state will be modified by an operation before it runs. Some services provide support to determine which objects are modified by an operation and to access their previous state, for example, relational databases. For other services, it is always possible to use coarse-grained objects and to call `modify` conservatively; this makes it simpler to write correct code at the expense of some performance degradation.

Non-determinism. BASE implements a form of state machine replication that requires replicas to behave deterministically. Our methodology uses abstraction to hide most of the nondeterminism in the implementations it reuses. However, many services involve forms of nondeterminism that cannot be hidden by abstraction. For instance, in the case of the NFS service, the time-last-modified for each file is set by reading the server's local clock. If this were done independently at each replica, the states of the replicas would diverge.

Instead, we allow the primary replica to propose values for non-deterministic choices by providing the `propose_value` upcall, which is only invoked at the primary. The call receives the client request and the sequence number for that request; it selects a non-deterministic value and puts it in `non-det`. This value is going to be supplied as an argument of the `execute` upcall to all replicas.

The protocol implemented by the BASE library prevents a faulty primary from causing replica state to diverge by sending different values to different backups. However, a faulty primary might send the same, incorrect value to all backups, subverting the system's desired behavior. The solution to this problem is to have each replica implement a `check_value` function that validates the choice of non-deterministic values that was made by the primary. If one-third or more non-faulty replicas reject a value proposed by a faulty primary, the request will not be executed and the view change mechanism will cause the primary to be replaced soon after.

Recovery. Proactive recovery periodically restarts each replica from a correct, up-to-date checkpoint of the abstract state that is obtained from the other replicas. Recoveries are triggered by a watchdog timer. When this timer fires, the replica reboots after saving to disk the abstract service state, and the replication protocol state, which includes abstract objects that were copied by the incremental checkpointing mechanism.

The library could invoke `get_obj` repeatedly to save a complete copy of the abstract state to disk but this would be expensive. It is sufficient to ensure that the current concrete state is on disk and to save a small amount of additional information to enable reconstruction of the conformance representation when the replica restarts. Since the library does not have access to this representation, the service state is saved to a file by an additional upcall, `shutdown`, that is implemented by the conformance wrapper. The conformance wrapper also implements a `restart` upcall that is invoked to reconstruct the conformance representation from the file saved by `shutdown` and from the concrete state of the service. This enables the replica to compute the abstract state by calling `get_obj` after `restart` completes.

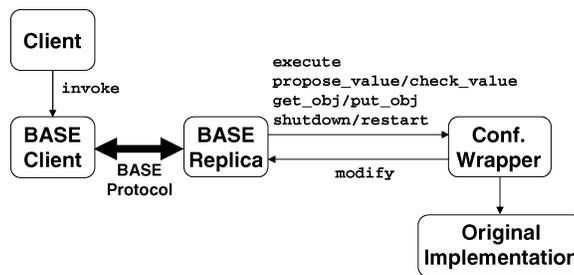


Fig. 2. BASE function calls and upcalls.

In some cases, the information in the conformance representation is volatile; it is no longer valid when the replica restarts. In this case, it is necessary to augment it with information that is persistent and allows restart to reconstruct the conformance representation after a reboot.

After calling `restart`, the library uses the hierarchical state transfer mechanism to compare the value of the abstract state of the replica with the abstract state values stored by the other replicas. It computes cryptographic hashes of the abstract objects and compares them with the hashes in the state partition tree to check if the objects are corrupt. The state partition tree also contains the sequence number of the last checkpoint when each object was modified [Castro and Liskov 2000]. The replica uses this information to check which objects are out-of-date without having to compute their hash. These checks are performed in parallel with fetches of objects that have already been determined to be out-of-date or corrupt. This is efficient: the replica fetches only the value of objects that are out-of-date or corrupt. We use a single threaded implementation with event queues representing objects to fetch and objects to check. Checks are performed while waiting for replies to fetch requests. The replica does not execute operations until it completes the recovery but the other replicas continue to process requests [Castro and Liskov 2000].

Proactive recovery allows automatic recovery of faulty replicas from many failures. However, hardware failures may prevent replicas from recovering automatically. In this case, administrators should repair or replace the faulty replica and then trigger the recovery mechanism to bring it up to date.

The object values fetched by the replica could be supplied to `put_objs` to update the concrete state, but the concrete state might still be corrupt. For example, an implementation may have a memory leak and simply calling `put_objs` will not free unreferenced memory. In fact, implementations will not typically offer an interface that can be used to fix all corrupt data structures in their concrete state. Therefore, it is better to restart the implementation from a clean initial concrete state and use the abstract state to bring it up-to-date. A global view of all BASE functions and upcalls that are invoked is shown in Figure 2.

2.4 Limitations

In theory, the methodology can be used to build a replicated service from any set of existing implementations of any service. But sometimes this may be hard because of the following problems.

Undocumented behavior. To apply the methodology, we need to understand and model the behavior of each service implementation. We do not need to model low level implementation details but only the behavior that can be observed by the clients of that implementation. We believe that the behavior of most software is well documented at this level, and we can use black box testing to understand small omissions in the documentation and small deviations from documented behavior. Implementations whose behavior we cannot model are unlikely to be of much use. It may be possible to remove operations whose behavior is not well documented from the abstract specification, or to implement these operations entirely in the conformance wrapper.

Very different behavior. If the implementations used to build the service behave very differently, any common abstract specification will deviate significantly from the behavior of some implementations. Theoretically, it is possible to write arbitrarily complex conformance wrappers and state conversion functions to bridge the gap between the behavior of the different implementations and the common abstract specification. In the worst case, we could implement the entire abstract specification in the wrapper code. But this is not practical because it is expensive to write complex wrappers, and complex wrappers are more likely to introduce new bugs. Therefore, it is important to use a set of implementations with similar behavior.

Narrow interfaces. The external interface of some implementations may not allow the wrapping code to read or write data that has an impact on the behavior observed by the client. For example, databases do not usually provide interfaces to manipulate concurrency control state, which influences observable behavior. There are three options in this case. First, the data can be shadowed in the conformance wrapper. This is practical if it is a small amount of data that is simple to maintain. Second, it may be possible to change the abstract specification such that this data has no impact on the behavior observed by the client. Third, it may be possible to gain access to internal APIs that avoid the problem.

Concurrency. It is common for service implementations to process requests concurrently to improve performance. Additionally, these implementations can provide different consistency guarantees for concurrent execution. For example, databases provide different degrees of isolation [Gray and Reuter 1993] to allow applications to trade off consistency for performance. Concurrent request execution results in a form of nondeterminism that is visible to the clients of the service and, therefore, needs to be constrained to apply our replication methodology (or any other form of state machine replication as discussed, for example, in Kemme and Alonso [2000]; Amir et al. [2002]; Jiménez-Peris et al. [2002]). It is nontrivial to ensure deterministic behavior without degrading performance. There are two basic approaches to solve this problem: implementing concurrency control in the conformance wrapper, or modifying the concurrency control code in the service implementation.

The conformance wrapper can implement concurrency control by determining which requests conflict and by not issuing a request to the service if it

conflicts with a request that has a smaller sequence number and has not yet completed. This approach treats the underlying service implementation as a black box and it works even when the service provides weak consistency guarantees for concurrent request execution or when different implementations provide different guarantees. However, determining conflicts before executing requests may be hard in some services. It is easy in the two examples presented in this article: file systems and client-server databases that ship transaction read and write sets back to the server. But it is harder in other systems, for example, relational databases where servers receive transactions with arbitrary sequences of SQL statements. The wrapper can conservatively assume that all requests conflict, which is simple to implement and solves the problem, but can result in poor performance. The work in Amir et al. [2002] and Jiménez-Peris et al. [2002] issues requests to a database one at a time to ensure determinism.

Alternatively, the concurrency control code in the service implementation can be modified to ensure that conflicting requests are serialized in order of increasing sequence number. This has the disadvantage of requiring nontrivial changes to the service implementation and it does not work with weak consistency guarantees. The work in Kemme and Alonso [2000] describes how to modify a relational database to achieve something similar.

3. EXAMPLES

This section uses two examples to illustrate the methodology: a replicated file system and an object oriented database.

3.1 File System

The file system is based on the NFS protocol [RFC-1094 1989]. Its replicas can run different operating systems and file system implementations. This allows them to tolerate software errors not only in the file system implementation but also in the rest of the operating system.

3.1.1 Abstract Specification. The common abstract specification is based on the specification of version 2 of the NFS protocol [RFC-1094 1989]. The abstract file service state consists of a fixed-size array of pairs containing an object and a generation number. Each object has a unique identifier, *oid*, which is obtained by concatenating its index in the array and its generation number. The generation number is incremented every time the entry is assigned to a new object. There are four types of objects:

- files, whose data is a byte array with the file contents;
- directories, whose data is a sequence of <name, oid> pairs ordered lexicographically by name;
- symbolic links, whose data is a small character string;
- special *null* objects, which indicate that an entry is free.

In addition to data, all non-null objects have meta-data, which includes the attributes in the NFS *fattr* structure, and the index (in the array) of its parent directory. Each entry in the array is encoded using XDR [RFC-1014 1987]. The

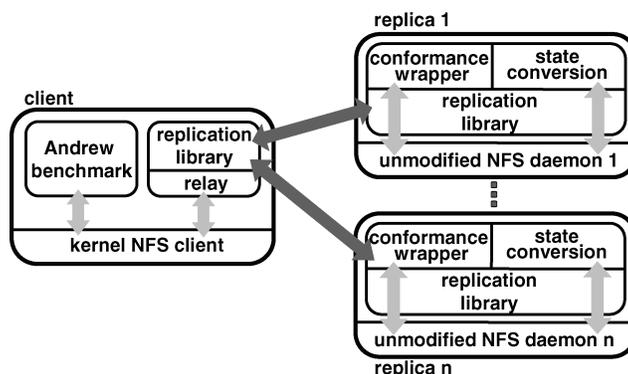


Fig. 3. Software architecture.

object with index 0 is a directory object that corresponds to the root of the file system tree that was mounted.

Keeping a pointer to the parent directory is redundant, since we can derive this information by scanning the rest of the abstract state. But it simplifies the inverse abstraction function and the recovery algorithm, as we will explain later.

The operations in the common specification are those defined by the NFS protocol. There are operations to read and write each type of non-null object. The file handles used by the clients are the *oids* of the corresponding objects. To ensure deterministic behavior, we require that *oids* be assigned deterministically, and that directory entries returned to a client be ordered lexicographically.

Some errors in the NFS protocol depend on the environment, for example, NFSERR_NOSPC is returned when the disk is full. The common abstract specification virtualizes the environment to ensure deterministic error processing. For example, the abstract state records the total number of bytes in abstract objects and a maximum capacity in bytes for the replicated file system. The abstract operations compare these values to decide when to raise NFSERR_NOSPC. The abstract state also records the maximum file size and name size to process NFSERR_FBIG and NFSERR_NAMETOOLONG deterministically. These maximum sizes and the disk capacity must be such that no correct concrete implementation raises the errors if they are not exceeded in the abstract state.

The abstraction hides many details: the allocation of file blocks, the representation of large files and directories, and the persistent storage medium and how it is accessed. This is desirable for simplicity and performance. Additionally, abstracting from implementation details like resource allocation improves resilience to software faults due to aging because proactive recovery can fix resource leaks.

3.1.2 Conformance Wrapper. There is a conformance wrapper around each implementation to ensure that it behaves according to the common specification. The conformance wrapper for the file service processes NFS protocol operations and interacts with an off-the-shelf file system implementation using the NFS protocol as illustrated in Figure 3. A file system exported by the replicated

file service is mounted on the client machine like any regular NFS file system. Application processes run unmodified and interact with the mounted file system through the NFS client in the kernel. We rely on user level relay processes to mediate communication between the standard NFS client and the replicas. A relay receives NFS protocol requests, calls the `invoke` procedure of our replication library, and sends the result back to the NFS client. The replication library invokes the `execute` procedure implemented by the conformance wrapper to run each NFS request. This architecture is similar to BFS [Castro and Liskov 1999].

The conformance representation consists of an array that corresponds to the one in the abstract state but it does not store copies of the objects; instead each array entry contains the type of object, the generation number, and for non-empty entries it also contains the file handle assigned to the object by the underlying NFS server, the value of the timestamps in the object's abstract meta-data, and the index of the parent directory. The abstract objects' data and remaining meta-data attributes are computed from the concrete state when necessary. The representation also contains a map from file handles to *oids* to aid in processing replies efficiently.

The wrapper processes each NFS request received from a client as follows. It translates the file handles in the request, which encode *oids*, into the corresponding NFS server file handles. Then it sends the modified request to the underlying NFS server. The server processes the request and returns a reply.

The wrapper parses the reply and updates the conformance representation. If the operation created a new object, the wrapper allocates a new entry in the array in the conformance representation, increments the generation number, and updates the entry to contain the file handle assigned to the object by the NFS server and the index of the parent directory. If any object is deleted, the wrapper marks its entry in the array free. In both cases, the reverse map from file handles to *oids* is updated.

The wrapper must also update the abstract timestamps in the array entries corresponding to objects that were accessed. For this, it uses the value for the current clock chosen by the primary using the `propose_value` upcall in order to prevent the states of the replicas from diverging. However, if a faulty primary chooses an incorrect value the system could have an incorrect behavior. For example, the primary might always propose the same value for the current time; this would cause all replicas to update the modification time to the same value that it previously held and therefore, according to the cache consistency protocol implemented by most NFS clients [Callaghan 1999], cause the clients to erroneously not invalidate their cached data, thus leading to inconsistent values at the caches of different clients. The solution to this problem is to have each replica validate the choice for the current timestamp using the `check_value` function. This function checks that the proposed timestamp is within a specified delta of the replica's own clock value, and that the timestamps produced by the primary are monotonically increasing. This always guarantees safety: all replicas agree on the timestamp of each operation, the timestamp is close to the clock reading of at least one correct replica, and timestamps are monotonically increasing. But we rely on loosely synchronized

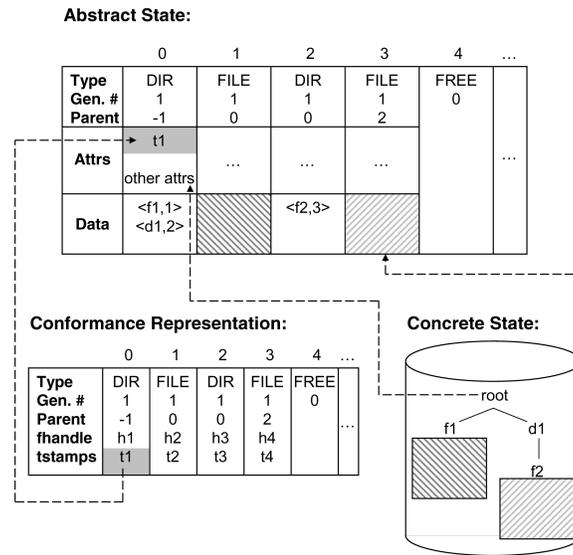


Fig. 4. Example of the abstraction function.

clocks for liveness, which is reasonable if replicas use an algorithm like NTP [Mills 1992].

Finally, the wrapper returns a modified reply to the client, using the map to translate file handles to *oids* and replacing the concrete timestamp values by the abstract ones. When handling *readdir* calls the wrapper reads the entire directory and sorts it lexicographically to ensure that the client receives identical replies from all replicas.

In the current implementation, the conformance wrapper issues read-write requests to the service one at a time to ensure that they are serialized in order of increasing sequence number at all the replicas. Read-only requests are processed using BFT's read-only optimization [Castro and Liskov 2002] and may be processed concurrently. We could improve performance by implementing a simple form of concurrency control in the wrapper and allowing non-conflicting read-write requests to execute concurrently.

This wrapper is simple and small, which is important because it reduces the likelihood of introducing additional software errors, and its implementation can be reused for all NFS server implementations.

3.1.3 State Conversions. The abstraction function in the file service is implemented as follows. For each file system object, it uses the file handle stored in the conformance representation to invoke the NFS server to obtain the data and meta-data for the object. Then it replaces the concrete timestamp values by the abstract ones, converts the file handles in directory entries to *oids*, and sorts the directories lexicographically.

Figure 4 shows how the concrete state and the conformance representation are combined to form the abstract state for a particular example. Note that the attributes in the concrete state are combined with the timestamps in the

```

function put_objs(in modified object array)
  for each entry  $\in$  modified object array do
    if new object's type = file or symbolic link then
      update_directory(new object's parent directory index)
      update object's meta-data in the conformance representation
      set file attributes
      write file's new contents or update link's target
    else if new object's type = directory then
      update_directory(this object's index)
      update object's meta-data in the conformance representation
    else if new object's type = free entry then
      update object's type in the conformance representation

function update_directory(in directory's index)
  if (directory has already been updated or
  directory has not changed) then
    do nothing
  else
    update_directory(new directory's parent directory index)
    read directory's old contents using NFS calls
    for each entry in old directory do
      if entry is not in new directory then
        remove entry using NFS call
      else if entry is in new directory but type is wrong then
        remove entry
    for each entry in new directory do
      if entry is not in old directory then
        create entry using NFS call

```

Fig. 5. Inverse abstraction function.

conformance representation to form the attributes in the abstract state. Also note that the contents of the files and directories are not stored by the conformance representation, but only in the concrete state.

The pseudo code for the inverse abstraction function in the file service is shown in Figure 5. This function receives an array with the indices of the objects that need to be updated and the new values for those objects. It scans each entry in the array to determine the type of the new object, and acts accordingly.

If the new object is a file or a symbolic link, it starts by calling the `update_directory` function, passing the new object's parent directory index as an argument. This will cause the object's parent directory to be reconstructed if needed, and the corresponding object in the underlying file system will be created if it did not exist already. Then it can update the object's entry in the conformance representation, and issue a `setattr` and a `write` to update the file's meta-data and data in the concrete state. For symbolic links, it is sufficient to update their meta-data.

When the new object is a directory, it is sufficient to invoke `update_directory` passing its own index as an argument, and then updating the appropriate entry in the conformance representation.

Finally, if the new object is a free entry it updates the conformance representation to reflect the new object's type and generation number. If the entry was not previously free, it must also remove the mapping from the file handle that

was stored in that entry to its *oid*. We do not have to update the parent directory of the old object, since it must have changed and will eventually be processed.

The `update_directory` function can be summarized as follows. If the directory that is being updated has already been updated or is not in the array of objects that need to be updated then the function performs no action. Otherwise it calls itself recursively passing the index of the parent directory (taken from the new object) as an argument. Then, it looks up the contents of the directory by issuing a `readdir` call. It scans the entries in the old state to remove the ones that are no longer present in the abstract state (or have a different type) and finally scans the entries in the new abstract state and creates the ones that are not present in the old state. When an entry is created or deleted, the conformance representation is updated to reflect this.

3.1.4 Proactive Recovery. After a recovery, a replica must be able to restore its abstract state. This could be done by saving the entire abstract state to disk before the recovery, but that would be very expensive. Instead we want to save only the metadata (e.g., the oids and the timestamps). But to do this we need a way of relating the oids to the files in the concrete file system state. This cannot be done using file handles since they can change when the NFS server restarts. However, the NFS specification states that each object is uniquely identified by a pair of meta-data attributes: $\langle \text{fsid}, \text{fileid} \rangle$. We solve the problem by adding another component to the conformance representation: a map from $\langle \text{fsid}, \text{fileid} \rangle$ pairs to the corresponding *oids*. The `shutdown` method saves this map (as well as the metadata maintained by the conformance representation for each file) to disk.

After rebooting, the `restart` method performs the following steps. It reads the map from disk; performs a new `mount` RPC call, thus obtaining the file handle for the file system root; and places null file handles in all the other entries in the conformance representation that correspond to all the other objects, indicating that we do not know the new file handles for those objects yet. It then initializes the other entries using the metadata that was stored by `shutdown`.

Then the replication library runs the protocol to bring the abstract state of the replica up to date. As part of this process, it updates the digests in its partition tree using information collected from the other replicas and calls `get_obj` on each object to check if it has the correct digest. This checks the integrity not only of file and directory contents but also of all their meta-data. Corrupt or out-of-date objects are fetched from the other replicas.

The call to `get_obj` determines the new NFS file handle if necessary. In this case, it goes up the directory tree (using the parent index in the conformance representation) until it finds a directory whose new file handle is already known. Then it issues a `readdir` to learn the names and fileids of the entries in the directory, followed by a `lookup` call for each one of those entries to obtain their NFS file handles; these handles are then stored in the array position that is determined by the $\langle \text{fsid}, \text{fileid} \rangle$ to *oid* map. Then it continues down the path of the object whose file handle is being reconstructed, computing not only the file handles of the directories in that path, but also those of all their siblings in the tree.

When walking up the directory tree using the parent indices, we need to detect loops so that the recovery function will not enter an infinite loop due to erroneous information stored by the replica during shutdown.

Currently, we restart the NFS server in the same file system and update its state with the objects fetched from other replicas. We could change the implementation to start an NFS server on a second empty disk and bring it up to date incrementally as we obtain the values of the abstract objects. This has the advantage of improving fault-tolerance as discussed in Section 2. Additionally, it can improve disk locality by clustering blocks from the same file and files that are in the same directory.

3.2 Object-Oriented Database

We have also applied our methodology to replicate the servers in the Thor object-oriented database [Liskov et al. 1999]. In this example, all the replicas run the same server implementation. The example is interesting because the service is more complex than NFS, and the server implementation is multithreaded and exhibits a significant degree of non-determinism. The methodology enabled reuse of the existing server code and could enable software rejuvenation through proactive recovery. We begin by giving a brief overview of Thor and then describe how the methodology was applied in this example. A more detailed description can be found in Rodrigues [2001].

3.2.1 System Overview. Thor [Liskov et al. 1999] provides a persistent object store that can be shared by applications running concurrently at different locations. It guarantees type-safe sharing by ensuring that all objects are used in accordance with their types. Additionally, it provides atomic transactions [Gray and Reuter 1993] to guarantee strong consistency in the presence of concurrent accesses and crashes.

Thor is implemented as a client/server system in which servers provide persistent storage for objects and applications run at clients on cached copies of persistent objects.

Servers store objects in *pages* on disk and also cache these pages in main memory. Each object stored by a server is identified by a 32-bit *oref*, which contains a *pagenum* that identifies the page where the object is stored and an *onum* that identifies the object within the page. Objects are uniquely identified by a pair containing the object *oref* and the identifier of the server that stores the object.

Each client maintains a cache of objects retrieved from servers in main memory [Castro et al. 1997]. Applications running at the client invoke methods on these cached objects. When an application requests an object that is not cached, the client fetches the page that contains the object from the corresponding server.

Thor uses an optimistic concurrency control algorithm [Adya et al. 1995] to *serialize* [Gray and Reuter 1993] transactions. Clients run transactions on cached copies of objects assuming that these copies are up-to-date but record *orefs* of objects read or modified by the transaction. To commit a transaction,

the client sends a commit request to the server that stores these objects. (Thor uses a two-phase commit protocol [Gray and Reuter 1993] when transactions access objects at multiple servers but we will not describe this case to simplify the presentation.) The commit request includes a transaction timestamp assigned by the client, the orefs it recorded, and the new values of modified objects.

The server attempts to serialize transactions in order of increasing timestamps. To determine if a transaction can commit, the server uses a *validation queue* (VQ) and *invalid sets* (ISs). The VQ contains an entry for every transaction that committed recently. Each entry contains the orefs of objects that were read or modified by the transaction, and the transaction's timestamp. There is an IS for each active client that lists orefs of objects with stale copies in that client's cache. A transaction can commit if none of the objects it accessed is in the corresponding IS, if it did not modify an object that was read by a committed transaction in the VQ with a later timestamp, and if it did not read an object that was modified by a transaction in the VQ with a later timestamp. If the transaction commits, its effects are recorded persistently; otherwise, it has no effect. In either case, the server informs the client of its decision.

The server updates the ISs of clients when a transaction commits. It adds orefs of objects modified by the transaction to the ISs of clients that are caching those objects. It computes this set of clients using a *cached-pages directory* that maps each page in the server to the set of clients that may have cached copies of that page. The server adds clients to the directory when they fetch pages and clients piggyback information about pages that they have discarded in fetch and commit requests that they send to the server. Similarly, the servers piggyback invalidation messages on fetch and commit replies to inform clients of objects in their IS. An object is removed from a client's IS when the server receives an acknowledgement for the invalidation. These acknowledgements are also piggybacked on other messages.

When a transaction commits, clients send new versions of modified objects but not their containing pages. These objects are stored by the server in a *modified object buffer* (MOB) [Ghemawat 1995] that allows the server to defer installing these objects to their pages on disk. The modifications are installed to disk lazily by a background flusher thread when the MOB is almost full to make room for new modifications.

3.2.2 Abstract Specification. We applied our methodology to replicate Thor servers. The abstract specification models the behavior of these servers as seen by clients. The interface exported by servers has four main operations: *start_session* and *end_session*, which are used by clients to start and end sessions with a server; and *fetch* and *commit*, which were described before. Invalidation messages are piggybacked on fetch and commit replies, and invalidation acknowledgements and notifications of page evictions from client caches are piggybacked on fetch and commit requests.

The abstract state of the service is defined as follows. The array of abstract objects is partitioned into four fixed-size areas.

Database pages. Each entry in this area corresponds to a page in the database. The value of the entry with index i is equal to the value of the page with pagenum i .

Validation queue. Entries in this area correspond to entries in the VQ. The value of each entry contains the timestamp that was assigned to the corresponding transaction (or a null timestamp for free entries), the status of the transaction, an array with the orefs of objects read by the transaction, and an array with the orefs of objects written by the transaction. When a transaction commits, it is assigned the free entry with the lowest index. When there are no free entries, the entry of the transaction with the lowest timestamp t is discarded to free an entry for a new transaction and any transaction that attempts to commit with timestamp lower than t is aborted [Liskov et al. 1999].

Note that entries are not ordered by timestamp because this could lead to inefficient checkpoint computation and state transfer. Inserting entries in the middle of an ordered sequence could require shifting a large number of entries. This would increase the cost of our incremental checkpointing technique and could increase the amount of data sent during state transfers.

Invalid sets. Each entry in this area corresponds to the invalid set of an active client. The value of an entry contains the client identifier (or a null identifier for free entries), and an array with the orefs of invalid objects. When a new client invokes *start_session*, it is assigned an *abstract client number* that corresponds to the index of its entry in this area. The entry is discarded when the client invokes *end_session*.

Cached-pages directory. There is one entry in this area per database page. The index of an entry is equal to the pagenum of the corresponding page minus the starting index for the area. The value of an entry is an array with the abstract numbers of clients that cache the page.

The abstraction hides the details of how the page cache and the MOB are managed at the servers. This allows different replicas to cache different pages, or install objects to disk pages at different times without having their abstract states diverge.

3.2.3 Conformance Wrapper. Thor servers illustrate one of the problems that make applying our methodology harder. The external interface they offer is too narrow to implement state conversion functions that are both simple and efficient. For example, the interface between clients and servers does not allow reading or writing the validation queue, the invalid sets, or the cached-pages directory.

We could solve this problem by shadowing this data in the wrapper but this is not practical because it would require reimplementing the concurrency control algorithm. Instead, we implemented the state conversion functions using internal APIs. This was possible because we had access to the server source code. We used these internal APIs as black boxes; we did not add new operations or change existing operations. These internal APIs were used only to implement the state conversion functions. They were not used to define the abstract

specification. This is important because we want this specification to abstract as many implementation details as possible.

We also replaced the communication library used between servers and clients by one with the same interface that calls the BASE library. This avoids the need for interposing client and server proxies, which was the technique we used in the file system example.

The conformance wrapper maintains only two data structures: the *VQ array* and the *client array*, which are used in the state conversion functions as we will describe next. Each entry in the VQ array corresponds to the entry with the same index in the VQ area of the abstract state, and it contains the transaction timestamp in that abstract entry. When a transaction commits, the wrapper assigns it an entry in the VQ array (as described in the abstract specification) and stores its timestamp there.

The entries in the client array are used to map abstract client numbers to the per-client data structures maintained by Thor. They are updated by the wrapper when clients start and end sessions with the server.

In Thor, transaction timestamps are assigned by clients. The conformance wrapper rejects timestamps that deviate more than a threshold from the time when the commit request is received. This is important to prevent faulty clients from committing transactions with very large timestamps, which could cause spurious aborts. The conformance wrapper uses the `propose_value` and `check_value` upcalls offered by the BASE library for replicas to agree on the time when the commit request is received. Replicas use the agreed upon value to decide whether to reject or accept the proposed timestamp. This ensures that all correct replicas reach the same decision.

Besides maintaining these two data structures and checking timestamps, the wrapper simply invokes the operations exported by the Thor server after calling `modify` to inform the BASE library of which abstract objects are about to be modified.

In the current implementation, the wrapper issues requests to the server one at a time to ensure that replicas agree on the fate of conflicting transactions in concurrent commit requests. We could improve performance by allowing more concurrency. For example, we could perform concurrency control checks and insert transactions into the VQ sequentially in order of increasing sequence number while allowing the rest of the execution to proceed concurrently.

3.2.4 State Conversions. The `get_obj` upcall receives the index of an abstract object and returns a pointer to a buffer containing the current value of that abstract object. The implementation of `get_obj` in this example uses the index to determine which area the abstract object belongs to. Then, it computes the value of the abstract object using the procedure that corresponds to the object's area:

Database pages. If the abstract object is a database page, `get_obj` retrieves a copy of the page from disk (or from the page cache) and applies any pending modifications to the page that are in the MOB. This is the current value of the page that is returned.

Validation queue. If the object represents a validation queue entry, `get_obj` retrieves the timestamp that corresponds to this entry from the VQ array in the conformance representation. Then it uses the timestamp to fetch the entry from the VQ maintained by the server, and copies the sets with orefs of objects read or modified by the transaction to compose the value of the abstract object.

Invalid sets. If the object represents an invalid set for a client with number c , `get_obj` uses the client array in the conformance representation to map c to the client data structure maintained by the server for the corresponding client. Then, it retrieves the client invalid set from this data structure and uses it to compose the abstract object value.

Cached-pages directory. In this case, `get_obj` determines the pagenum of the requested abstract object by computing the offset to the beginning of the area. Then, it uses the pagenum to look up the information to compose the abstract object value in the cached-pages directory maintained by the server.

The `put_objs` upcall receives an array with new values for abstract objects and updates the concrete state to match these values. It iterates over the abstract object values and uses the object indices to determine which of the procedures below to execute.

Database pages. To update a concrete database page, `put_objs` removes any modifications in the MOB for that page to ensure that the new page value will not be overwritten with old modifications. Then, it places a page matching the new abstract value in the server's cache and marks it as dirty.

Validation queue, invalid sets and cached-pages directory. If the relevant server data structure already contains an entry corresponding to a new abstract object value, the function just updates the entry according to the new value. Otherwise, it must delete the entry from the server data structure if the new abstract object value describes a nonexistent entry, or create the entry if it did not previously exist and fill in the values according to the new abstract value. The conformance representation is updated accordingly.

4. EVALUATION

Our replication technique must achieve two goals to be successful: it must have low overhead, and the code of the conformance wrapper and the state conversion functions must be simple. It is important for the code to be simple to reduce the likelihood of introducing more errors and to keep the monetary cost of using our technique low. This section evaluates the extent to which both example applications meet each of these goals.

A detailed performance evaluation of the BFT library appears in Castro and Liskov [2000, 2002] and Castro [2000]. It includes several micro-benchmarks to evaluate performance in the normal case, during view changes, and during state transfer. Additionally, it includes experiments to evaluate the performance impact of the various optimizations used by BFT. Furthermore, these experimental results are backed by an analytic performance model [Castro 2000]. These results are relevant to the evaluation of the BASE library. The difference

is that BASE adds the cost of conversions between abstract and concrete representations.

4.1 File System Overhead

This section presents results of experiments that compare the performance of our replicated file system with the off-the-shelf, unreplicated NFS implementations that it wraps.

4.1.1 Experimental Setup. Our technique has three advantages: reuse of existing code, software rejuvenation using proactive recovery, and opportunistic N-version programming. We present results of experiments to measure the overhead in systems that benefit from different combinations of these advantages. We ran experiments with and without proactive recovery in a *homogeneous* setup, where all replicas ran the same operating system, and in a *heterogeneous* setup, where each replica ran a different operating system.

All experiments ran with four replicas and one client. Four replicas can tolerate one Byzantine fault; we expect this reliability level to suffice for most applications. Experiments to evaluate the performance of the replication algorithm with more clients and replicas appear in Castro [2000].

In the homogeneous setup, clients and replicas ran on Dell Precision 410 workstations with Linux 2.2.16-3 (uniprocessor). These workstations have a 600 MHz Pentium III processor, 512 MB of memory, and a Quantum Atlas 10K 18WLS disk. All machines were connected by a 100 Mb/s switched Ethernet and had 3Com 3C905B interface cards. The switch was an Extreme Networks Summit48 V4.1. The experiments ran on an isolated network.

The heterogeneous setup used the same hardware setup but some replicas ran different operating systems. The client and one of the replicas ran Linux as in the homogeneous setup. The other replicas ran different operating systems: one ran Solaris 8 1/01; another ran OpenBSD 2.8; and the last one ran FreeBSD 4.0.

All experiments ran the modified Andrew benchmark [Howard et al. 1988; Ousterhout 1990], which emulates a software development workload. It has five phases: (1) creates subdirectories recursively; (2) copies a source tree; (3) examines the status of all the files in the tree without examining their data; (4) examines every byte of data in all the files; and (5) compiles and links the files. They ran the scaled up version of the benchmark described in Castro and Liskov [2000] where phase 1 and 2 create n copies of the source tree, and the other phases operate in all these copies. We ran a version of Andrew with n equal to 100, Andrew100, that creates approximately 200 MB of data and another with n equal to 500, Andrew500, that creates approximately 1 GB of data. Andrew100 fits in memory at both the client and the replicas but Andrew500 does not.

The benchmark ran at the client machine using the standard NFS client implementation in the Linux kernel with the following mount options: UDP transport, 4096-byte read and write buffers, allowing write-back client caching, and allowing attribute caching. All the experiments report the average of three

Table I. Andrew100: Elapsed Time in Seconds

Phase	BASEFS	NFS-std
1	0.9	0.5
2	49.2	27.4
3	45.4	39.2
4	44.7	36.5
5	287.3	234.7
Total	427.65	338.3

Table II. Andrew500: Elapsed Time in Seconds

Phase	BASEFS	NFS-std
1	5.0	2.4
2	248.2	137.6
3	231.5	199.2
4	298.5	238.1
5	1545.5	1247.1
Total	2328.7	1824.4

runs of the benchmark and the standard deviation was always below 7% of the reported values.

4.1.2 *Homogeneous Results.* Tables I and II present the results for Andrew100 and Andrew500 in the homogeneous setup with no proactive recovery. They compare the performance of our replicated file system, BASEFS, with the standard, unreplicated NFS implementation in Linux with Ext2fs at the server, NFS-std. In these experiments, BASEFS is also implemented on top of a Linux NFS server with Ext2fs at each replica.

The results show that the overhead introduced by our replication technique is low: BASEFS takes only 26% longer than NFS-std to run Andrew100 and 28% longer to run Andrew500. The overhead is different for the different phases mostly due to variations in the amount of time the client spends computing between issuing NFS requests.

There are two main sources of overhead: the cost of running the Byzantine-fault-tolerant replication protocol, and the cost of abstraction. The latter includes the time spent running the conformance wrapper and the time spent running the abstraction function to compute checkpoints of the abstract file system state. We estimate that the Byzantine fault tolerance protocol adds approximately 15% of overhead relative to NFS-std in Andrew100 and 20% in Andrew500. This estimate is based on the overhead of BFS relative to NO-REP for Andrew100 and Andrew500 that was reported in Castro and Liskov [2000]. We expect this estimate to be fairly accurate: BFS is very similar to BASEFS except that it does not use abstraction, and NO-REP is identical to BFS except that it is not replicated. The remaining overhead of 11% relative to NFS-std in Andrew100 and 8% in Andrew500 can be attributed to abstraction.

We also ran Andrew100 and Andrew500 with proactive recovery. The results, which are labeled BASEFS-PR, are shown in Table III. The results for Andrew100 were obtained by recovering replicas round robin with a

Table III. Andrew with Proactive Recovery:
Elapsed Time to Run the Benchmark in Seconds

System	Andrew100	Andrew500
BASEFS-PR	448.2	2385.1
BASEFS	427.65	2328.7
NFS-std	338.33	1824.4

Table IV. Andrew: Maximum Time to
Complete a Recovery in Seconds

	Andrew100	Andrew500
Shutdown	0.07	0.32
Reboot	30.05	30.05
Restart	0.18	0.97
Fetch and check	18.28	141.37
Total	48.58	172.71

new recovery starting every 80 seconds; reboots were simulated by sleeping 30 seconds.¹ We obtained the results for Andrew500 in the same way but in this case a new recovery was started every 250 seconds. This leads to a window of vulnerability of approximately 6 minutes for Andrew100 and 17 minutes for Andrew500; that is, the system will work correctly as long as fewer than 1/3 of the replicas fail in a correlated way within any time window of size 6 (or 17) minutes. (A discussion on windows of vulnerability with proactive recovery appears in [Castro and Liskov 2002].) The results show that even with these very strong guarantees BASEFS is only 32% slower than NFS-std in Andrew100 and 31% slower in Andrew500.

Table IV presents a breakdown of the time to complete the slowest recovery in Andrew100 and Andrew500. *Shutdown* accounts for the time to write the state of the replication library and the conformance representation to disk, and *restart* is the time to read this information back. *Fetch and check* is the time to rebuild the *oid* to file handle mappings in the conformance wrapper, to convert the state stored by the NFS server to its abstract form and check it, and to fetch out-of-date objects from other replicas. Fetching out-of-date objects is done in parallel with converting and checking the state.

The recovery time in Andrew100 is dominated by the time to reboot but as the state size increases, reading, converting, and checking the state become the dominant cost; this accounts for 141 seconds in Andrew500 (82% of the total recovery time). Scaling to larger states is an issue but we could use the techniques suggested in Castro and Liskov [2000] that make the cost of checking proportional to the number of objects modified in a time period rather than to the total number of objects in the state.

As mentioned, we would like our implementation of proactive recovery to start an NFS server on a second empty disk with a clean file system to improve the range of faults that can be tolerated. We believe that extending our implementation in this way should not significantly affect the performance of

¹This reboot time is based on the results obtained by the LinuxBIOS project [Minnich 2000]. They claim to be able to reboot Linux in 35 s by replacing the BIOS with Linux.

Table V. Andrew100
Heterogeneous: Elapsed Time
in Seconds

System	Elapsed Time
BASEFS-PR	1950.6
BASEFS	1662.2
OpenBSD	1599.1
Solaris	1009.2
FreeBSD	848.4
Linux	338.3

the recovery. We would write each abstract object to the new file system asynchronously right after checking it. Since the value of the abstract object is already in memory at this point and it is written to a different disk, the additional overhead should be minimal.

4.1.3 Heterogeneous Results. Table V presents results for Andrew100 with and without proactive recovery in the heterogeneous setup. In this experiment, each BASEFS replica runs a different operating system with a different NFS and file system implementation. The table also presents results for the standard NFS implementation in each operating system without replication.

The overhead of BASEFS in this experiment varies from 4% relative to the slowest replica (OpenBSD) to 391% relative to the fastest replica (Linux). The replica running Linux is much faster than all the others because Linux does not ensure stability of modified data and meta-data before replying to the client as required by the NFS protocol. The overhead relative to OpenBSD is low because BASEFS only requires a quorum with 3 replicas, which must include the primary, to complete operations. These results were obtained with the primary replica in the machine running Linux. Therefore, BASEFS does not need to wait for the slowest replica to complete operations. However, this replica slows down the others because it is out-of-date most of the time and it is constantly transferring state from the others. This partially explains why the overhead relative to the third fastest replica (Solaris) is higher than in the homogeneous case (65% versus 26%).

We also ran BASEFS with proactive recoveries in the heterogeneous setup. We recovered a new replica every 425 seconds; reboots were simulated by sleeping 30 seconds. In this case, the overhead varies from 22% relative to the slowest replica to 477% relative to the fastest replica.

The overhead of BASEFS-PR relative to BASEFS without proactive recovery is higher in the heterogeneous setup than in the homogeneous setup. This happens because proactive recovery causes the slowest replica to periodically become the primary. During these periods the system must wait for the slowest replica to complete operations (and to get up-to-date before it can complete operations).

4.2 Object-Oriented Database Overhead

This section presents results of experiments to measure the overhead of our replicated implementation of Thor relative to the original implementation of

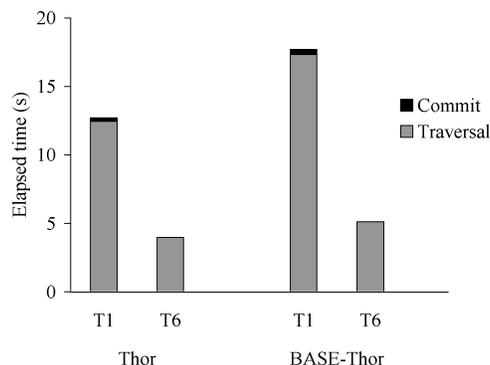


Fig. 6. OO7: Elapsed time, cold read-only traversals.

Thor without replication. To provide a conservative measurement, the version of Thor without replication does not ensure stability of information committed by a transaction. A real implementation would save a transaction log to disk or use replication to ensure stability as we do. In either case, the overhead introduced by BASE would be lower.

4.2.1 Experimental Setup. We ran four replicas and one client in the homogeneous setup described in Section 4.1. The experiments ran the OO7 benchmark [Carey et al. 1993], which is intended to match the characteristics of many different CAD/CAM/CASE applications. The OO7 database contains a tree of assembly objects, with leaves pointing to three composite parts chosen randomly from among 500 such objects. Each composite part contains a graph of atomic parts linked by connection objects; each atomic part has 3 outgoing connections. All our experiments ran on the medium database, which has 200 atomic parts per composite part.

The OO7 benchmark defines several database traversals; these perform a depth-first traversal of the assembly tree and execute an operation on the composite parts referenced by the leaves of this tree. Traversals T1 and T6 are read-only; T1 performs a depth-first traversal of the entire composite part graph, while T6 reads only its root atomic part. Traversals T2a and T2b are identical to T1 except that T2a modifies the root atomic part of the graph, while T2b modifies all the atomic parts. We ran each traversal in a single transaction.

The objects are clustered into 4 KB pages in the database. The database takes up 38 MB in our implementation. Each server replica had a 20 MB cache (of which 16 MB were used for the MOB); the client cache had 16MB. All the results we report are for cold traversals: the client and server caches were empty in the beginning of the traversals.

4.2.2 OO7 Results. The results in Figure 6 are for read-only traversals. We measured elapsed times for T1 and T6 traversals of the database, both in the original implementation, Thor, and the version that is replicated with BASE, BASE-Thor. The figure shows the total time to run the transaction broken into the time to run the traversal and the time to commit the transaction.

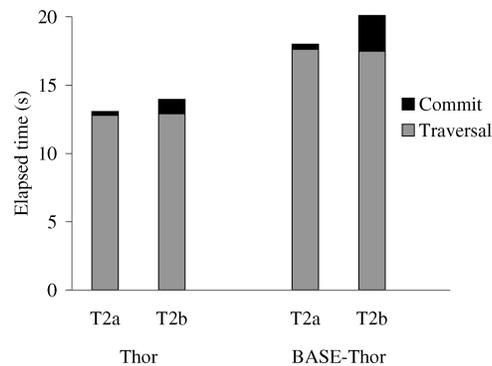


Fig. 7. Elapsed time, cold read-write traversals.

BASE-Thor takes 39% more time to complete T1, and 29% more time to complete T6. The commit cost is a small fraction of the total time in these experiments. Therefore, most of the overhead is due to an increase in the cost to fetch pages. The micro-benchmarks in Castro and Liskov [2000] predict an overhead of 60% when fetching 4 KB pages with no computation at the client or the replicas. The overhead here is lower because the pages have to be read from the replicas' disks. Similarly, the relative overhead is lower for traversal T6 because it generates disk accesses with less locality. Thus, the average time to read a disk page from the server disk is higher in T6 than in T1. We expect a similar effect in more realistic settings where the database does not fit in main memory at either the server or the clients. In these settings, BASE will have lower overhead because the cost of disk accesses will dominate performance.

Figure 7 shows elapsed times for read-write traversals. In this case, BASE adds an overhead relative to the original implementation of 38% in T2a and 45% in T2b. The traversal times for T1, T2a, and T2b are almost identical because these traversals are very similar. What is different is the time to commit the transactions. Traversal T2a modifies 500 atomic parts whereas T2b modifies 100000. Therefore, the commit time is a significant fraction of the total time in traversal T2b but not in traversal T2a. BASE increases the commit overhead significantly due to the cost of maintaining checkpoints. The overhead for read-write traversals would be significantly lower relative to a version of Thor that ensured stability of transaction logs.

4.3 Code Complexity

To implement the conformance wrapper and the state conversion functions, it is necessary to write new code. It is important for this code to be simple so that it is easy to write and not likely to introduce new bugs. We measured the number of semicolons in the code we wrote for the replicated file system and for the replicated database to evaluate its complexity. Counting semicolons is better than counting lines because it does not count comment and blank lines.

The code we wrote for the replicated file system has a total of 1105 semicolons with 624 in the conformance wrapper and 481 in the state conversion functions. Of the semicolons in the state conversion functions, only 45 are specific

to proactive recovery. The code of the conformance wrapper is trivial. It has 624 semi-colons only because there are many NFS operations. The code of the state conversion functions is slightly more complex because it involves directory tree traversals with several special cases but it is still rather simple. To put these numbers in perspective, the number of semicolons in the code in the Linux 2.2.16 kernel that is directly related with the file system, NFS, and the driver of our SCSI adapter is 17735. Furthermore, this represents only a small fraction of the total size of the operating system; Linux 2.2.6 has 544442 semicolons including drivers and 229095 semicolons without drivers.

The code we wrote for the replicated database has a total of 658 semicolons with 345 in the conformance wrapper and 313 in the state conversion functions. To put these numbers in perspective, the number of semicolons in the original Thor code is 37055.

5. RELATED WORK

The BASE library is implemented as an extension to the BFT library [Castro and Liskov 1999, 2000, 2002] and BASEFS is inspired by the BFS file system [Castro and Liskov 1999]. But BFT requires all replicas to run the same service implementation and does not allow reuse of existing code without significant modifications. Our technique for software rejuvenation [Huang et al. 1995] is based on the proactive recovery technique implemented in BFT [Castro and Liskov 2000]. But the use of abstraction allows us to tolerate software errors due to aging that could not be tolerated in BFT, for example, resource leaks in the service code. Additionally, it allows us to combine proactive recovery with N-version Programming.

There is a lengthy discussion of work related to BFT in Castro and Liskov [2002], which includes work not only on Byzantine fault tolerance but also on replication in general. Therefore, we omit a discussion of that work from this paper and concentrate on work related to what is new in BASE relative to BFT.

N-Version Programming [Chen and Avizienis 1978] exploits design diversity to reduce common mode failures. It works as follows: N software development teams produce different implementations of the same service specification for the same customer; the different implementations are then run in parallel; and voting is used to produce a common result. This technique has been criticized for several reasons [Gray and Siewiorek 1991]: it increases development and maintenance costs by a factor of N or more, and it adds unacceptable time delays to the implementation. In general, this is considered to be a powerful technique, but with limited usability since only a small subset of applications can afford it.

BASE enables low cost N-version programming by reusing existing implementations from different vendors. Since each implementation is developed for a large number of customers, there are significant economies of scale that keep the development, testing, and maintenance costs per customer low. Additionally, the cost of writing the conformance wrappers and state conversion functions is kept low by taking advantage of existing interoperability standards. The end result is that our technique will cost less and may actually be more

effective at reducing common mode failures because competitive pressures will keep implementations from different vendors independent.

Recovery of faulty versions has been addressed in the context of N-Version Programming [Romanovsky 2000; Tso and Avizienis 1987] but these approaches have suffered from two problems. First, they are inefficient and cannot scale to services with large state. Second, they require detailed knowledge of each version, which precludes our opportunistic N-Version programming technique. For example, Romanovsky [2000] proposes a technique where each version defines a conversion function from its concrete state to an abstract state. But this abstract state is based on what is common across the implementations of the different versions. Our technique improves on this by providing a very efficient recovery mechanism and by deriving the abstract state from an abstract behavioral specification that captures what is visible to the client succinctly; this leads to better fault tolerance and efficiency.

A different way to achieve better fault tolerance through diversity appears in Ramkumar and Strumpfen [1997]. This work provides a compiler-assisted mechanism for portable checkpointing that enables recovery of a service in a different machine with a different processor architecture or even a different operating system. This technique uses a single implementation but exploits diversity in an environment using checkpointing/recovery techniques.

Several other systems have used wrapping techniques to replicate existing components [Cooper 1985; Liskov et al. 1991; Bressoud and Schneider 1995; Maffeis 1995; Moser et al. 1998; Narasimhan et al. 1999]. Many of these systems have relied on standards like NFS or CORBA [Object Management Group 1999] to simplify wrapping of existing implementations. For example, Eternal [Moser et al. 1998] is a commercial implementation of the new Fault Tolerant CORBA standard [Object Management Group 2000]. All of these systems except Immune [Narasimhan et al. 1999] assume benign faults.

There are significant differences between these systems and BASE. First, they assume that replicas run identical implementations. They also assume that replicas are deterministic or they resolve the nondeterminism at a low level of abstraction. For example, many resolve nondeterminism by having a primary run the operations and ship the resulting state or a log with all nondeterministic events to the backups (e.g., Bressoud and Schneider [1995]). Not only does this fail to work with Byzantine faults, but replicas are more likely to fail at the same time because they are forced to behave identically at a low level of abstraction.

RNFS [Marzullo and Schmuck 1988] implements a replicated NFS file system from an existing implementation of NFS, and Postgres-R [Kemme and Alonso 2000] and the work in Amir et al. [2002] and Liménez-Peris et al. [2002] implement replicated databases from an existing implementation of Postgres [Stonebraker et al. 1990]. They use group communication toolkits like Isis [Birman et al. 1991] and Ensemble [Hayden 1998] to coordinate wrappers, and wrappers hide observable nondeterminism from the clients rather than forcing deterministic behavior at a low level. They differ from our work because they assume that all replicas run the same implementation, they cannot tolerate Byzantine faults, and they do not provide a proactive recovery mechanism.

BASE uses abstraction to hide most nondeterminism and to enable replicas to run different implementations. It also offers an efficient mechanism for replicas to agree on nondeterministic choices that works with Byzantine faults. This is important when these choices are directly visible by clients, for example, timestamps. Additionally, we provide application-independent support for efficient state transfer and for incremental conversion between abstract and concrete state, which is important because these are harder with Byzantine faults.

The work described in Salles et al. [1999] uses wrappers to ensure that an implementation satisfies an abstract specification. These wrappers use the specification to check the correctness of outputs generated by the implementation and contain faults. They are not used to enable replication with different or nondeterministic implementations as in BASE.

Since both the examples described in this article relate to storage, it is worthwhile comparing them with RAID [Chen et al. 1994], which is the most widely deployed replicated storage solution. RAID is implemented in hardware in many mother boards and it is cheap. However, BASEFS and BASE-Thor replicate not only the disk but the entire service implementation and operating system. They are more expensive but they may be able to tolerate faults that RAID cannot. For example, errors in the operating system in one replica could cause file system state to become corrupt. BASEFS may tolerate this error if it does not occur in the operating systems of other replicas, but any RAID solution would simply write corrupt information to the replicated disk.

6. CONCLUSION

Software errors are a major cause of outages and they are increasingly exploited in malicious attacks to gain control or deny access to important services. Byzantine fault tolerance allows replicated systems to mask some software errors but it has been expensive to deploy. We have described a replication technique, BASE, which uses abstraction to reduce the cost of deploying Byzantine fault tolerance and to improve its ability to withstand attacks and mask software errors.

BASE reduces cost because it enables reuse of off-the-shelf service implementations, and it improves resilience to software errors by enabling opportunistic N-version programming, and software rejuvenation through proactive recovery.

Opportunistic N-version programming runs distinct, off-the-shelf implementations at each replica to reduce the probability of common mode failures. To apply this technique, it is necessary to define a common abstract behavioral specification for the service and to implement appropriate conversion functions for the state, requests, and replies of each implementation in order to make it behave according to the common specification. These tasks are greatly simplified by basing the common specification on standards for the interoperability of software from different vendors; these standards appear to be common, for example, ODBC [Geiger 1995], and NFS [RFC-1094 1989]. Opportunistic N-version programming improves on previous N-version programming techniques by avoiding the high development, testing, and maintenance costs without compromising the quality of individual versions.

Additionally, we provide a mechanism to repair faulty replicas. Proactive recovery allows the system to remain available provided no more than one-third of the replicas become faulty and corrupt the abstract state (in a correlated way) within a window of vulnerability. Abstraction may enable service availability even when more than one-third of the replicas are faulty because it can hide corrupt items in concrete states of faulty replicas.

The paper described BASEFS—a replicated NFS file system implemented using our technique. The conformance wrapper and the state conversion functions in our prototype are simple, which suggests that they are unlikely to introduce new bugs and that the monetary cost of using our technique would be low.

We ran the Andrew benchmark to compare the performance of our replicated file system and the off-the-shelf implementations that it reuses. Our performance results indicate that the overhead introduced by our technique is low; BASEFS performs within 32% of the standard NFS implementations that it reuses.

We also used the methodology to build a Byzantine fault-tolerant version of the Thor object-oriented database [Liskov et al. 1999] and made similar observations. In this case, the methodology enabled reuse of the existing database code, which is nondeterministic.

As future work, it would be interesting to apply the BASE technique to a relational database service by taking advantage of the ODBC standard. Additionally, a library of mappings between abstract and concrete states for common data structures would further simplify our technique.

ACKNOWLEDGMENTS

We would like to thank Chandrasekhar Boyapati, João Garcia, Ant Rowstron, and Larry Peterson for their helpful comments on drafts of this paper. We also thank Charles Blake, Benjie Chen, Dorothy Curtis, Frank Dabek, Michael Ernst, Kevin Fu, Frans Kaashoek, David Mazières, and Robert Morris for help in providing an infrastructure to run some of the experiments.

REFERENCES

- ADYA, A., GRUBER, R., LISKOV, B., AND MAHESHWARI, U. 1995. Efficient Optimistic Concurrency Control using Loosely Synchronized Clocks. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*. San Jose, California, 23–34.
- AMIR, Y., DANILOV, C., MISKIN-AMIR, M., STANTON, J., AND TUTU, C. 2002. Practical Wide-Area Database Replication. Tech. Rep. CNDS-2002-1, Johns Hopkins University.
- BIRMAN, K., SCHIPER, A., AND STEPHENSON, P. 1991. Lightweight causal and atomic group multicast. In *ACM Trans. Comput. Syst.* Vol. 9(3).
- BRESSOUD, T. AND SCHNEIDER, F. 1995. Hypervisor-based Fault Tolerance. In *Proceedings of the Fifteenth ACM Symposium on Operating System Principles*. Copper Mountain Resort, Colorado, 1–11.
- CALLAGHAN, B. 1999. *NFS Illustrated*. Addison-Wesley, Reading, Massachusetts.
- CAREY, M. J., DEWITT, D. J., AND NAUGHTON, J. F. 1993. The OO7 Benchmark. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*. Washington D.C., 12–21.

- CASTRO, M. 2000. Practical Byzantine fault-tolerance. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- CASTRO, M., ADYA, A., LISKOV, B., AND MYERS, A. 1997. HAC: Hybrid Adaptive Caching for Distributed Storage Systems. In *Proceedings of the Sixteenth ACM Symposium on Operating System Principles*. Saint Malo, France, 102–115.
- CASTRO, M. AND LISKOV, B. 1999. Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*. New Orleans, Louisiana, 173–186.
- CASTRO, M. AND LISKOV, B. 2000. Proactive recovery in a Byzantine-fault-tolerant system. In *Proceedings of the Fourth Symposium on Operating Systems Design and Implementation*. San Diego, California, 273–288.
- CASTRO, M. AND LISKOV, B. 2002. Practical Byzantine Fault Tolerance and Proactive Recovery. *ACM Trans. Comput. Syst.* 20, 4 (Nov.), 398–461.
- CHEN, L. AND AVIZIENIS, A. 1978. N-Version Programming: A Fault-Tolerance Approach to Reliability of Software Operation. In *Fault Tolerant Computing, FTCS-8*. 3–9.
- CHEN, P. M., LEE, E. K., GIBSON, G. A., KATZ, R. H., AND PATTERSON, D. A. 1994. RAID: High-performance, reliable secondary storage. *ACM Comput. Surv.* 26, 2, 145–185.
- COOPER, E. 1985. Replicated Distributed Programs. In *Proceedings of the Tenth ACM Symposium on Operating System Principles*. Orcas Island, Washington, 63–78.
- GEIGER, K. 1995. *Inside ODBC*. Microsoft Press.
- GHEMAWAT, S. 1995. The Modified Object Buffer: a storage management technique for object-oriented databases. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- GRAY, J. AND SIEWIOREK, D. 1991. High-availability computer systems. *IEEE Comput.* 24, 9 (Sept.), 39–48.
- GRAY, J. N. AND REUTER, A. 1993. *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann Publishers Inc.
- HAYDEN, M. 1998. The Ensemble System. Tech. Rep. TR98-1662, Cornell University, Ithaca, New York. Jan.
- HERLIHY, M. P. AND WING, J. M. 1987. Axioms for Concurrent Objects. In *Conference Record of the Fourteenth Annual ACM Symposium on Principles of Programming Languages*. Munich, Germany, 13–26.
- HOWARD, J., KAZAR, M., MENEES, S., NICHOLS, D., SATYANARAYANAN, M., SIDEBOTHAM, R., AND WEST, M. 1988. Scale and performance in a distributed file system. *ACM Trans. Comput. Syst.* 6, 1 (Feb.), 51–81.
- HUANG, Y., KINTALA, C., KOLETTIS, N., AND FULTON, N. D. 1995. Software rejuvenation: Analysis, modules and applications. In *Digest of Papers: FTCS-25, The Twenty-Fifth International Symposium on Fault-Tolerant Computing*. Pasadena, California, 381–390.
- JIMÉNEZ-PERIS, R., PATIÑO-MARTÍNEZ, M., KEMME, B., AND ALONSO, G. 2002. Improving the Scalability of Fault-Tolerant Database Clusters. In *Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS'02)*. Vienna, Austria.
- KEMME, B. AND ALONSO, G. 2000. Don't be lazy be consistent: Postgres-R, a new way to implement Database Replication. In *VLDB 2000, Proceedings of 26th International Conference on Very Large Data Bases*. Cairo, Egypt.
- LAMPORT, L. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Comm. ACM* 21, 7 (July), 558–565.
- LISKOV, B., CASTRO, M., SHRIRA, L., AND ADYA, A. 1999. Providing persistent objects in distributed systems. In *Proceedings of the 13th European Conference on Object-Oriented Programming (ECOOP '99)*. Lisbon, Portugal, 230–257.
- LISKOV, B., GHEMAWAT, S., GRUBER, R., JOHNSON, P., SHRIRA, L., AND WILLIAMS, M. 1991. Replication in the Harp File System. In *Proceedings of the Thirteenth ACM Symposium on Operating System Principles*. Pacific Grove, California, 226–238.
- LISKOV, B. AND GUTTAG, J. 2000. *Program Development in Java: Abstraction, Specification, and Object-Oriented Design*. Addison-Wesley.
- MAFFEIS, S. 1995. Adding group communication and fault tolerance to CORBA. In *Proceedings of the Second USENIX Conference on Object-Oriented Technologies*. Toronto, Canada, 135–146.

- MARZULLO, K. AND SCHMUCK, F. 1988. Supplying high availability with a standard network file system. In *Proceedings of the 8th International Conference on Distributed Computing Systems*. San Jose, California, 447–453.
- MILLS, D. L. 1992. Network Time Protocol (Version 3) Specification, Implementation and Analysis. Network Working Report RFC 1305.
- MINNICH, R. 2000. The Linux BIOS Home Page. <http://www.acl.lanl.gov/linuxbios>.
- MOSER, L., MELLIAR-SMITH, P., AND NARASIMHAN, P. 1998. Consistent object replication in the eternal system. *Theory and Practice of Object Systems* 4, 2 (Jan.), 81–92.
- NARASIMHAN, P., KIHSTROM, K., MOSER, L., AND MELLIAR-SMITH, P. 1999. Providing Support for Survivable CORBA Applications with the Immune System. In *Proceedings of the 19th International Conference on Distributed Computing Systems*. Austin, Texas, 507–516.
- OBJECT MANAGEMENT GROUP. 1999. The Common Object Request Broker: Architecture and Specification. OMG technical committee document formal/98-12-01. June.
- OBJECT MANAGEMENT GROUP. 2000. Fault Tolerant CORBA. OMG technical committee document orbos/2000-04-04. Mar.
- OUSTERHOUT, J. 1990. Why Aren't Operating Systems Getting Faster as Fast as Hardware? In *Proceedings of the Usenix Summer 1990 Technical Conference*. Anaheim, California, 247–256.
- PEASE, M., SHOSTAK, R., AND LAMPORT, L. 1980. Reaching Agreement in the Presence of Faults. *J. ACM* 27, 2 (Apr.), 228–234.
- RAMKUMAR, B. AND STRUMPEN, V. 1997. Portable checkpointing for heterogeneous architectures. In *Digest of Papers: FTCS-27, The Twenty-Seventh Annual International Symposium on Fault-Tolerant Computing*. Seattle, Washington, 58–67.
- RFC-1014 1987. Network working group request for comments: 1014. XDR: External data representation standard.
- RFC-1094 1989. Network working group request for comments: 1094. NFS: Network file system protocol specification.
- RODRIGUES, R. 2001. Combining abstraction with Byzantine fault-tolerance. M.S. thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- ROMANOVSKY, A. 2000. Faulty version recovery in object-oriented N-version programming. *IEEE Proc. Soft.* 147, 3 (June), 81–90.
- SALLES, F., RODRÍGUEZ, M., FABRE, J., AND ARLAT, J. 1999. MetaKernels and Fault Containment Wrappers. In *Digest of Papers: FTCS-29, The Twenty-Ninth Annual International Symposium on Fault-Tolerant Computing*. Madison, Wisconsin, 22–29.
- SCHNEIDER, F. 1990. Implementing fault-tolerant services using the state machine approach: a tutorial. *ACM Comput. Surv.* 22, 4 (Dec.), 299–319.
- STONEBRAKER, M., ROWE, L. A., AND HIROHAMA, M. 1990. The implementation of POSTGRES. *IEEE Trans. Knowl. Data Eng.* 2, 1 (Mar.), 125–142.
- TSO, K. AND AVIZIENIS, A. 1987. Community error recovery in N-version software: A design study with experimentation. In *Digest of Papers: FTCS-17, the Seventeenth Annual Symposium on Fault Tolerant Computing*. Pittsburgh, Pennsylvania, 127–133.
- W3C. 2000. Extensible Markup Language (XML) 1.0 (Second Edition). W3C recommendation.

Received September 2001; revised November 2002; accepted December 2002