

A Cooperative File System

by

Frank Dabek

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2001

© Massachusetts Institute of Technology 2001. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
5 September, 2001

Certified by.....
M. Frans Kaashoek
Professor of Computer Science and Engineering
Thesis Supervisor

Certified by.....
Robert T. Morris
Assistant Professor of Computer Science and Engineering
Thesis Supervisor

Accepted by.....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

A Cooperative File System

by
Frank Dabek

Submitted to the Department of Electrical Engineering and Computer Science
on 5 September, 2001, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Computer Science and Engineering

Abstract

The Cooperative File System (CFS) is a new peer-to-peer read-only storage system that provides provable guarantees for the efficiency, robustness, and load-balance of file storage and retrieval. CFS does this with a completely decentralized architecture that can scale to large systems. CFS servers provide a distributed hash table (DHash) for block storage. CFS clients interpret DHash blocks as a file system. DHash distributes and caches blocks at a fine granularity to achieve load balance, uses replication for robustness, and decreases latency with server selection. DHash finds blocks using the Chord location protocol, which operates in time logarithmic in the number of servers and requires logarithmic state at each node.

CFS is implemented using the SFS file system toolkit and runs on many UNIX operating systems including Linux, OpenBSD, and FreeBSD. Experience on a globally deployed prototype shows that CFS delivers data to clients as fast as FTP. Controlled tests show that CFS is able to route queries in a scalable way. For example, in experiments with a system of 4,096 servers, looking up a block of data involves contacting only seven servers. In general, a logarithmic number of servers must be contacted to route a query. Servers are also able to join and leave the system efficiently. Tests demonstrate nearly perfect robustness and unimpaired performance even when as many as half the servers fail.

Thesis Supervisor: M. Frans Kaashoek

Title: Professor of Computer Science and Engineering

Thesis Supervisor: Robert T. Morris

Title: Assistant Professor of Computer Science and Engineering

Acknowledgments

This thesis is the result of joint work with Frans Kaashoek, Robert Morris, David Karger and Ion Stoica. Most of what is presented here is also described in “Wide Area Storage with CFS” published in the proceedings of SOSP 2001.

I am especially indebted to Frans for his advice, support, and, above and beyond the call of adviser, his implementation of the Chord protocol on which this work is based. Also to Robert, for providing the initial impetus and continuing guidance for “Plan Three” which eventually became CFS and for his assistance with the presentation of this thesis.

This work would never have come to fruition without the assistance of David Mazières who answered countless questions about his very helpful user-level file system framework and who taught me more than I ever wanted to know about C++. David and Kevin Fu’s work on SFSRO also heavily influenced the design of CFS.

I acknowledge Chuck Blake for the helpful advice he provided and the thought-provoking alternative suggestions he put forth regarding the design of CFS.

All of the members of PDOS provided a supportive working atmosphere. Thanks especially to Eric Peterson who was good company on more than one late night paper marathon.

The Defense Advanced Research Projects Agency (DARPA) provided funding for this work under contract N66001-00-1-8933.

Finally, to Mary, who provided the sustenance that allowed me to see this project to completion and, not only that, bakes a mean oatmeal raisin cookie.

Contents

1	Introduction	9
1.1	CFS: A Cooperative File System	9
1.2	Contributions	10
1.3	Thesis Overview	11
2	Related Work	13
2.1	Naming and Authentication	13
2.2	Peer-to-Peer Search	13
2.3	Anonymous Storage	14
2.4	Peer-to-Peer Hash Based Systems	14
2.5	Web Caches	15
3	Design Overview	17
3.1	System Structure	17
3.2	Design Motivation	18
3.3	CFS Properties	20
4	Chord Layer	21
4.1	Consistent Hashing	21
4.2	The Chord Lookup Algorithm	21
4.3	Node ID Authentication	23
5	DHash Layer	25
5.1	Replication	25
5.2	Caching	26
5.3	Load Balance	27
5.4	Quotas	28
5.5	Updates and Deletion	28
6	Server Selection	29
6.1	Correctness of server selection	30
7	Implementation	33
7.1	Chord Implementation	33
7.2	DHash Implementation	33
7.3	Client Implementation	35
7.4	Using CFS	37

8	Experimental Results	39
8.1	Real Life	39
8.2	Controlled Experiments	42
8.2.1	Server Selection	42
8.2.2	Lookup Cost	43
8.2.3	Load Balance	43
8.2.4	Caching	44
8.2.5	Storage Space Control	47
8.2.6	Effect of Failure	47
9	Future Work and Conclusions	51

Chapter 1

Introduction

Existing peer-to-peer systems (such as Napster [24], Gnutella [12], and Freenet [6]) demonstrate the benefits of cooperative storage and serving. These systems have the potential to be more available and tolerate faults better than their client-server counterparts. In addition, peer-to-peer systems offer the unique advantage of being able to harness idle network and storage resources of large numbers of participating servers.

Achieving these benefits requires overcoming a number of design challenges. A peer-to-peer architecture should be decentralized and symmetric to prevent a small group of nodes from limiting overall system performance under high load. Similarly, the system should not present a single point of failure. Symmetry also implies that load should be well balanced across the available servers in proportion to their ability or willingness to handle load.

We expect that participants in a peer-to-peer system will be volunteers using hardware and networks that are less well managed and less reliable than a typical server deployed today. A peer-to-peer system should operate well under these conditions. In particular, the system should support join and leave operations for servers efficiently as we expect that participants will be active for relatively short periods of time, on the order of one hour.

Finding desired data in a large system must be efficient; file transfers from a peer-to-peer storage system should be competitive with point-to-point transfers using the FTP or HTTP protocols.

Finally, the system should maintain these properties even as it grows very large. Reports place the peak number of registered users of the Napster system at approximately 70 million and the number of active users at 1.5 million [7]. Any viable design should be able to support a user base on this order of magnitude in size.

1.1 CFS: A Cooperative File System

While the peer-to-peer systems in common use solve some of these problems, none solves all of them. This thesis presents CFS (the Cooperative File System), a new design that meets all of these challenges. Two major components of CFS are presented here in detail: a distributed block store we call DHash and the CFS file system itself which utilizes the DHash block store. The Chord distributed lookup system is used by DHASH and provides many of the system's desirable properties; Chord is described in detail in [35]. Figure 1.1 summarizes the CFS software layering:

A CFS file system exists as a set of blocks distributed over the available CFS servers by DHash. CFS client software interprets the stored blocks as file system data and meta-data

Layer	Responsibility
FS	Interprets blocks as files; presents a file system interface to applications.
DHash	Stores unstructured data blocks reliably.
Chord	Maintains routing tables used to find blocks.

and presents an ordinary read-only file-system interface to applications. The client that we have implemented and present here uses an NFS loop back server [19] to access stored data as a file system mapped into the local namespace. It is not the only client we could have chosen to implement; one could use the flexible DHash block store to implement, for example, a system similar in spirit to Napster but without a centralized database.

The DHash (distributed hash) layer stores and retrieves unstructured data blocks for the client, distributes the blocks among the servers and maintains cached and replicated copies. DHash uses Chord to locate the servers responsible for a block. DHash layers block management on top of the Chord distributed hash function. DHash provides load balance for popular large files by arranging to spread the blocks of each file over many servers. To balance the load imposed by popular small files, DHash caches each block at servers likely to be consulted by future Chord lookups for that block. DHash supports asynchronous operation to allow for pre-fetching to decrease download latency. DHash replicates each block at a small number of servers, to provide fault tolerance. DHash enforces weak quotas on the amount of data each server can inject, to deter abuse. Finally, DHash allows control over the number of *virtual servers* per server, to provide control over how much data a server must store on behalf of others.

The Chord lookup system implements a hash-like operation that maps from block identifiers to servers. Chord assigns each server an identifier drawn from the same 160-bit identifier space as block identifiers. These identifiers can be thought of as points on a circle. The mapping that Chord implements takes a block's ID and yields the block's *successor*, the server whose ID most closely follows the block's ID on the identifier circle. To implement this mapping, Chord maintains at each server a table with information about $O(\log N)$ other servers, where N is the total number of servers. A Chord lookup sends messages to $O(\log N)$ servers to consult their tables. Using Chord, CFS can find data efficiently even with a large number of servers, and servers can join and leave the system with few table updates. To improve Chord's performance we have modified it to route queries preferentially to hosts nearby in the underlying network topology. We present initial results of those experiments here.

1.2 Contributions

CFS has been implemented. This thesis reports experimental results from a small international deployment of CFS servers and from a large-scale controlled test-bed. These results confirm the contributions of the CFS design:

- an aggressive approach to load balance by spreading file blocks randomly over servers;

- download performance on an Internet-wide prototype deployment as fast as standard FTP;
- provable efficiency and provably fast recovery times after failure;
- network proximity aware routing;
- simple algorithms to achieve the above results.

CFS is not yet in operational use, and such use will likely prompt refinements to its design. One potential area for improvement is the ability of the Chord lookup algorithm to tolerate malicious participants by verifying the routing information received from other servers. Another area that CFS does not currently address is anonymity; it is expected that anonymity, if needed, would be implemented as a new software layer above CFS, or perhaps by modifying the DHash system.

1.3 Thesis Overview

This thesis begins with a survey of related work in the field of peer-to-peer storage and lookup (Section 2). Section 3 provides an overview of the CFS system and explains how the Chord, DHash and CFS software layers interact. Section 4 presents an overview of the Chord lookup system and Section 5 discusses the DHash layer in detail. In Section 6 we discuss an extension to the Chord system developed for this work: server selection. Section 7 describes the implementation of CFS and presents pseudocode for its core algorithms. The results of experiments with an implementation of CFS are presented in Section 8. Finally, we present conclusions and potential future work in Section 9.

Chapter 2

Related Work

CFS was inspired by Napster [24], Gnutella [12], and particularly Freenet [6]. CFS uses peer-to-peer distributed hashing similar in spirit to a number of ongoing research projects [29, 32, 39]. In comparison to existing peer-to-peer file sharing systems, CFS offers simplicity of implementation and high performance without compromising correctness. CFS balances server load, finds data quickly, and guarantees data availability in the face of server failures with very high probability. CFS, as a complete system, has individual aspects in common with many existing systems. The major relationships are summarized below.

2.1 Naming and Authentication

CFS authenticates data by naming it with public keys or content hashes, as do many other distributed storage systems [10, 6, 8, 15, 32, 38]. The use of content-hashes to securely link together different pieces of data is due to Merkle [21]; the use of public keys to authentically name data is due to the SFS system [20].

CFS adopts naming, authentication, and file system structure ideas from SFSRO [10], which implements a secure distributed read-only file system—that is, a file system in which files can be modified only by their owner, and only through complete replacement of the file. However, SFSRO and CFS have significant differences at the architectural and mechanism levels. SFSRO defines protocols and authentication mechanisms which a client can use to retrieve data from a given server. CFS adds the ability to *dynamically* find the server currently holding the desired data, via the Chord location service. This increases the robustness and the availability of CFS, since changes in the set of servers are transparent to clients.

2.2 Peer-to-Peer Search

Napster [24] and Gnutella [12] are arguably the most well known peer-to-peer file systems today. They present a keyword search interface to clients, rather than retrieving uniquely identified data. As a result they are more like search engines than distributed hash tables, and they trade scalability for this power: Gnutella broadcasts search queries to many machines, and Napster performs searches at a central facility. CFS as described in this paper doesn't provide search, but we are developing a scalable distributed search engine for CFS.

Mojo Nation [22] is a broadcast query peer-to-peer storage system which divides files into blocks and uses a secret sharing algorithm to distribute the blocks to a number of hosts.

CFS also divides files into blocks but does not use secret sharing.

The Morpheus system [23] works analogously to Napster but selects a fraction of participating peers to serve as indexing servers, eliminating the single point of failure present in Napster. Morpheus allows for the striping of downloads from multiple servers, but unlike CFS it stores copies of whole files at each server. Both Napster and Morpheus do not explicitly replicate documents, but rather depend on the fact that popular documents will be widely downloaded and republished by users.

2.3 Anonymous Storage

Freenet [5] uses probabilistic routing to preserve the anonymity of clients, publishers, and servers. This anonymity requirement limits Freenet's reliability and performance. Freenet avoids associating a document with any predictable server, and avoids forming any globally coherent topology among servers. The former means that unpopular documents may simply disappear from the system, since no server has the responsibility for maintaining replicas. The latter means that a search may need to visit a large fraction of the Freenet network. As an example, Hong shows in his Figure 14-12 that in a network with 1,000 servers, the lookup path length can exceed 90 hops [26]. This means that if the hop count is limited to 90, a lookup may fail even though the document is available. Because CFS does not try to provide anonymity, it can guarantee much tighter bounds on lookup cost; for example, in a 4,096-node system, lookups essentially never exceed 10 hops.

CFS's caching scheme is similar to Freenet's in the sense that both leave cached copies of data along the query path from client to where the data was found. Because CFS finds data in significantly fewer hops than Freenet, and CFS's structured lookup paths are more likely to overlap than Freenet's, CFS can make better use of a given quantity of cache space.

Like Freenet, Publius [38] focuses on anonymity, but achieves it with encryption and secret sharing rather than routing. Publius requires a static, globally-known list of servers; it stores each share at a fixed location that is predictable from the file name. Free Haven [8] uses both cryptography and routing (using re-mailers [4]) to provide anonymity; like Gnutella, Free Haven finds data with a global search.

CFS does not attempt to provide anonymity, focusing instead on efficiency and robustness. We believe that intertwining anonymity with the basic data lookup mechanism interferes with correctness and performance. On the other hand, given a robust location and storage layer, anonymous client access to CFS could be provided by separate anonymizing proxies, using techniques similar to those proposed by Chaum [4] or Reiter and Rubin [30].

2.4 Peer-to-Peer Hash Based Systems

CFS layers storage on top of an efficient distributed hash lookup algorithm. A number of recent peer-to-peer systems use approaches similar to chord and offer similar scalability and performance, including CAN [29], PAST [31, 32], OceanStore [15, 39], and Ohaha [25]. A detailed comparison of these algorithms can be found in [35].

The PAST [32] storage system differs from CFS in its approach to load balance. Because a PAST server stores whole files, a server might not have enough disk space to store a large file even though the system as a whole has sufficient free space. A PAST server solves this by offloading files it is responsible for to servers that do have spare disk space. PAST handles the load of serving popular files by caching them along the lookup path.

CFS stores blocks, rather than whole files, and spreads blocks evenly over the available servers; this prevents large files from causing unbalanced use of storage. CFS solves the related problem of different servers having different amounts of storage space with a mechanism called *virtual servers*, which gives server managers control over disk space consumption. CFS's block storage granularity helps it handle the load of serving popular large files, since the serving load is spread over many servers along with the blocks. This is more space-efficient, for large files, than whole-file caching. CFS relies on caching only for files small enough that distributing blocks is not effective. Evaluating the performance impact of block storage granularity is one of the purposes of this thesis.

OceanStore [15] aims to build a global persistent storage utility. It provides data privacy, allows client updates, and guarantees durable storage. However, these features come at a price: complexity. For example, OceanStore uses a Byzantine agreement protocol for conflict resolution, and a complex protocol based on Plaxton trees [27] to implement the location service [39]. OceanStore assumes that the core system will be maintained by commercial providers.

LH* [17] is a distributed data structure based on linear hashing which supports constant time insertion and lookup of objects. Data stored in LH* is distributed to nodes via a split operation which occurs when a node reaches its capacity; splitting allows LH* to maintain a high utilization (a property not addressed by Chord which assumes an excess of available storage). LH*, however, assumes either that each node has knowledge of all other nodes in the system or that a central coordinator manages this information. This property is in conflict with the symmetry and scalability goals of Chord.

Ohaha [25] uses consistent hashing to map files and keyword queries to servers, and a Freenet-like routing algorithm to locate files. As a result, it shares some of the same weaknesses as Freenet.

2.5 Web Caches

Content distribution networks (CDNs), such as Akamai [1], handle high demand for data by distributing replicas on multiple servers. CDNs are typically managed by a central entity, while CFS is built from resources shared and owned by a cooperative group of users.

There are several proposed scalable cooperative Web caches [3, 9, 11, 18]. To locate data, these systems either multicast queries or require that some or all servers know about all other servers. As a result, none of the proposed methods is both highly scalable and robust. In addition, load balance is hard to achieve as the content of each cache depends heavily on the query pattern.

Cache Resolver [34], like CFS, uses consistent hashing to *evenly* map stored data among the servers [13, 16]. However, Cache Resolver assumes that clients know the entire set of servers; maintaining an up-to-date server list is likely to be difficult in a large peer-to-peer system where servers join and depart at unpredictable times.

Chapter 3

Design Overview

CFS provides distributed read-only file storage. It is structured as a collection of servers that provide block-level storage. Publishers (producers of data) and clients (consumers of data) layer file system semantics on top of this block store much as an ordinary file system is layered on top of a disk. Many unrelated publishers may store separate file systems on a single CFS system; the CFS design is intended to support the possibility of a single world-wide system consisting of millions of servers.

3.1 System Structure

Figure 3-1 illustrates the structure of the CFS software. Each CFS client contains three software layers: a file system client, a DHash storage layer, and a Chord lookup layer. The client file system uses the DHash layer to retrieve blocks. The client DHash layer uses the client Chord layer to locate the servers that hold desired blocks.

Each CFS server has two software layers: a DHash storage layer and a Chord layer. The server DHash layer is responsible for storing keyed blocks, maintaining proper levels of replication as servers come and go, and caching popular blocks. The server DHash and Chord layers interact in order to integrate looking up a block identifier with checking for cached copies of the block. CFS servers are oblivious to file system semantics: they simply provide a distributed block store.

CFS clients interpret DHash blocks in a file system format adopted from SFSRO [10]; the format is similar to that of the UNIX V7 file system, but uses DHash blocks and block identifiers in place of disk blocks and disk addresses. As shown in Figure 3-2, each block is




Figure 3-1: CFS software structure. Vertical links are local APIs; horizontal links are RPC APIs.




Figure 3-2: A simple CFS file system structure example. The root-block is identified by a public key and signed by the corresponding private key. The other blocks are identified by cryptographic hashes of their contents.

either a piece of a file or a piece of file system meta-data, such as a directory. The maximum size of any block is on the order of tens of kilobytes. A parent block contains the identifiers of its children.

The publisher inserts the file system’s blocks into the CFS system, using a hash of each block’s content (a content hash) as its identifier. Then the publisher signs the root block with his or her private key, and inserts the root block into CFS using the corresponding public key as the root block’s identifier. Clients name a file system using the public key; they can check the integrity of the root block using that key, and the integrity of blocks lower in the tree with the content hash identifiers that refer to those blocks. This approach guarantees that clients see an authentic and internally consistent view of each file system, though under some circumstances a client may see an old, but still internally consistent, version of a recently updated file system.

A CFS file system is read-only as far as clients are concerned. However, a file system may be updated by its publisher. This involves updating the file system’s root block in place, to make it point to the new data. CFS authenticates updates to root blocks by checking that the new block is signed by the same key as the old block. A timestamp prevents replays of old updates. CFS allows file systems to be updated without changing the root block’s identifier so that external references to data need not be changed when the data is updated.

CFS stores data for an agreed-upon finite interval. Publishers that want indefinite storage periods can periodically ask CFS for an extension; otherwise, a CFS server may discard data whose guaranteed period has expired. CFS has no explicit delete operation: instead, a publisher can simply stop asking for extensions. In this area, as in its replication and caching policies, CFS relies on the assumption that large amounts of spare disk space are available.

3.2 Design Motivation

CFS is intended to provide a highly available file system distributed across a very large collection of unreliable nodes. The system’s two main design decisions reflect this: CFS uses the Chord location algorithm which maintains only a small amount of state and allows efficient join and leave operations, and CFS distributes a single file over many hosts to balance load.

Splitting each file system (and file) into blocks and distributing those blocks over many servers balances the load of serving popular files over many servers. It also increases the number of messages required to fetch a whole file, since a client must look up each block separately. However, the network bandwidth consumed by a lookup is small compared to the bandwidth required to deliver the block. In addition, CFS hides the block lookup

latency by pre-fetching blocks.

Systems such as Freenet [6] and PAST [32] store whole files. This results in lower lookup costs than CFS, one lookup per file rather than per block, but requires more work to achieve load balance. Servers unlucky enough to be responsible for storing very large files may run out of disk space even though the system as a whole has sufficient free space. Balancing the load of serving whole files typically involves adaptive caching. Again, this may be awkward for large files; a popular file must be stored in its entirety at each caching server. DHash also uses caching, but only depends on it for small files.

DHash's block granularity is particularly well suited to serving large, popular files, such as software distributions. For example, in a 1,000-server system, a file as small as 8 megabytes will produce a reasonably balanced serving load with 8 KByte blocks; the load of serving the file will be distributed across 1,000 servers. A system that balances load by caching whole files would require, in this case, about 1,000 times as much total storage to achieve the same load balance since the whole file would be required to be cached on 1,000 servers. On the other hand, DHash is not as efficient as a whole-file scheme for large but unpopular files, though the experiments in Section 8.1 show that it can provide competitive download speeds. DHash's block granularity is not likely to affect (for better or worse) performance or load balance for small files. For such files, DHash depends on caching and on server selection among block replicas (described in Section 5.1).

CFS's core location algorithm, Chord, requires a logarithmic number of messages per lookup and provides a logarithmic bound on the number of messages a join or leave operation requires. CFS might have used the more naive approach of simply maintaining a list of every node on each node and then using consistent hashing or some other distributed storage algorithm to store and retrieve data. This approach would allow for constant time lookups but requires a number of messages linear in the number of nodes in the system for each join or leave operation; additionally this scheme requires storage linear in the number of nodes at each node. While an approach that requires linear state could be implemented for systems with a small collection of nodes or a larger system where the set of participating nodes rarely changed, it is not appropriate for the environment we intend CFS to be used in: millions of unreliable machines. A measurement study of Napster and Gnutella hosts reports that 50 percent of sessions lasted less than 60 minutes and that 25 percent lasted less than approximately 10 minutes [33].

In an environment with such a large number of join and leave operations, we feel that Chord offers an appealing tradeoff between the cost of performing a lookup and the cost of a membership change. Maintaining an accurate group membership over a wide-area network is difficult. Keidar et al. propose an algorithm to maintain group membership in a WAN but only demonstrate scaling to hundreds of clients and utilize a central server which CFS explicitly attempts to avoid [14]. Additionally, any large cache of node locations is likely to contain entries which have become unreachable (a linear group membership list could not be easily maintained). These unreachable entries will not affect the correctness of a lookup but will cause the lookup to suffer a large number of expensive network timeouts.

A detailed description and analysis of the additional mechanisms Chord uses to achieve load balance and reliability are presented in Section 5.

3.3 CFS Properties

CFS provides consistency and integrity of file systems by adopting the SFSRO file system format. CFS extends SFSRO by providing the following desirable properties:

- **Decentralized control.** CFS servers need not share any administrative relationship with publishers. CFS servers could be ordinary Internet hosts whose owners volunteer spare storage and network resources.
- **Scalability.** CFS lookup operations use space and messages at most logarithmic in the number of servers.
- **Availability.** A client can always retrieve data as long as it is not trapped in a small partition of the underlying network, and as long as one of the data's replicas is reachable using the underlying network. This is true even if servers are constantly joining and leaving the CFS system. CFS places replicas on servers likely to be at unrelated network locations to ensure independent failure.
- **Load balance.** CFS ensures that the burden of storing and serving data is divided among the servers in rough proportion to their capacity. It maintains load balance, even if some data are far more popular than others, by caching blocks and by spreading each file's data over many servers.
- **Persistence.** Once CFS commits to storing data, it keeps it available for at least an agreed-on interval.
- **Quotas.** CFS limits the amount of data that any particular IP address can insert into the system. This provides a degree of protection against malicious attempts to exhaust the system's storage.
- **Efficiency.** Clients can fetch CFS data with delay comparable to that of FTP, due to CFS's use of efficient lookup algorithms, caching, pre-fetching, and server selection.

The next two sections present Chord and DHash, which together provide these properties.

Chapter 4

Chord Layer

CFS uses the Chord protocol to locate blocks [35]. Chord supports just one operation: given a key, it will determine the node responsible for that key. Chord does not itself store keys and values, but provides primitives that allow higher-layer software to build a wide variety of storage systems; CFS is one such use of the Chord primitive. This section summarizes Chord; for a more detailed description of Chord see one of the several publications relating to Chord [35] [36].

4.1 Consistent Hashing

Each Chord node has a unique m -bit node identifier (ID), obtained by hashing the node's IP address and a *virtual node index*. Chord views the IDs as occupying a circular identifier space. Keys are also mapped into this ID space, by hashing them to m -bit key IDs. Chord defines the node responsible for a key to be the *successor* of that key's ID. The *successor* of an ID j is the node with the smallest ID that is greater than or equal to j (with wrap-around), much as in consistent hashing [13].

Consistent hashing lets nodes enter and leave the network with minimal movement of keys. To maintain correct successor mappings when a node n joins the network, certain keys previously assigned to n 's successor become assigned to n . When node n leaves the network, all of n 's assigned keys are reassigned to its successor. No other changes in the assignment of keys to nodes need occur.

Consistent hashing is straightforward to implement, with constant-time lookups, if all nodes have an up-to-date list of all other nodes. However, such a system does not scale; Chord provides a scalable, distributed version of consistent hashing.

4.2 The Chord Lookup Algorithm

A Chord node uses two data structures to perform lookups: a successor list and a finger table. Only the successor list is required for correctness, so Chord is careful to maintain its accuracy. The finger table accelerates lookups, but does not need to be accurate, so Chord is less aggressive about maintaining it. The following discussion first describes how to perform correct (but slow) lookups with the successor list, and then describes how to accelerate them up with the finger table. This discussion assumes that there are no malicious participants in the Chord protocol; while we believe that it should be possible for nodes to verify the

routing information that other Chord participants send them, the algorithms to do so are left for future work.

Every Chord node maintains a list of the identities and IP addresses of its r immediate successors on the Chord ring. The fact that every node knows its own successor means that a node can always process a lookup correctly: if the desired key is between the node and its successor, the latter node is the key's successor; otherwise the lookup can be forwarded to the successor, which moves the lookup strictly closer to its destination.

A new node n learns of its successors when it first joins the Chord ring, by asking an existing node to perform a lookup for n 's successor; n then asks that successor for its successor list. The r entries in the list provide fault tolerance: if a node's immediate successor does not respond, the node can substitute the second entry in its successor list. All r successors would have to simultaneously fail in order to disrupt the Chord ring, an event that can be made very improbable with modest values of r . An implementation should use a fixed r , chosen to be $2 \log_2 N$ for the foreseeable maximum number of nodes N .

The main complexity involved with successor lists is in notifying an existing node when a new node should be its successor. The stabilization procedure described in [35] does this in a way that guarantees to preserve the connectivity of the Chord ring's successor pointers.

Lookups performed only with successor lists would require an average of $N/2$ message exchanges, where N is the number of servers. To reduce the number of messages required to $O(\log N)$, each node maintains a finger table with m entries. The i^{th} entry in the table at node n contains the identity of the *first* node that succeeds n by at least 2^{i-1} on the ID circle. Thus every node knows the identities of nodes at power-of-two intervals on the ID circle from its own position. A new node initializes its finger table by querying an existing node. Existing nodes whose finger table or successor list entries should refer to the new node find out about it by periodic lookups performed as part of an asynchronous, ongoing stabilization process.

Figure 4-1 shows pseudo-code to look up the successor of identifier id . The main loop is in *find_predecessor*, which sends *preceding_node_list* RPCs to a succession of other nodes; each RPC searches the tables of the other node for nodes yet closer to id . Each iteration will set n' to a node between the current n' and id . Since *preceding_node_list* never returns an ID greater than id , this process will never overshoot the correct successor. It may under-shoot, especially if a new node has recently joined with an ID just before id ; in that case the check for $id \notin (n', n'.successor]$ ensures that *find_predecessor* persists until it finds a pair of nodes that straddle id .

Two aspects of the lookup algorithm make it robust. First, an RPC to *preceding_node_list* on node n returns a list of nodes that n believes are between it and the desired id . Any one of them can be used to make progress towards the successor of id ; they must all be unresponsive for a lookup to fail. Second, the **while** loop ensures that *find_predecessor* will keep trying as long as it can find any next node closer to id . As long as nodes are careful to maintain correct successor pointers, *find_predecessor* will eventually succeed.

In the usual case in which most nodes have correct finger table information, each iteration of the **while** loop eliminates half the remaining distance to the target. This means that the hops early in a lookup travel long distances in the ID space, and later hops travel small distances. The efficacy of the caching mechanism described in Section 5.2 depends on this observation. It is worthwhile to note that this algorithm does not itself provide an $O(\lg N)$ bound; the structure of the finger table, which the algorithm examines, guarantees that each hop will cover half of the remaining distance. This behavior is the source of the algorithm's logarithmic properties.

```

// Ask node n to find id's successor; first
// finds id's predecessor, then asks that
// predecessor for its own successor.
n.find_successor(id)
  n' = find_predecessor(id);
  return n'.successor();

// Ask node n to find id's predecessor.
n.find_predecessor(id)
  n' = n;
  while (id ∉ (n', n'.successor()))
    l = n'.preceding_node_list(id);
    n' = max n'' ∈ l s.t. n'' is alive
  return n';

// Ask node n for a list of nodes in its finger table or
// successor list that precede id.
n.preceding_node_list(id)
  return {n' ∈ {fingers ∪ successors}
         s.t. n' ∈ (n, id]}

```

Figure 4-1: The pseudo-code to find the successor node of an identifier id . Remote procedure calls are preceded by the remote node.

The following two theorems, proved in an accompanying technical report [36], show that neither the success nor the performance of Chord lookups is likely to be affected even by massive simultaneous failures. Both theorems assume that the successor list has length $r = O(\log N)$. A Chord ring is *stable* if every node's successor list is correct.

Theorem 1 *In a network that is initially stable, if every node then fails with probability $1/2$, then with high probability `find_successor` returns the closest living successor to the query key.*

Theorem 2 *In a network that is initially stable, if every node then fails with probability $1/2$, then the expected time to execute `find_successor` is $O(\log N)$.*

The evaluation in Section 8 validates these theorems experimentally.

4.3 Node ID Authentication

If Chord nodes could use arbitrary IDs, an attacker could destroy chosen data by choosing a node ID just after the data's ID. With control of the successor, the attacker's node could effectively delete the block by denying that the block existed.

To limit the opportunity for this attack, a Chord node ID must be of the form $h(x)$, where h is the SHA-1 hash function and x is the node's IP address concatenated with a virtual node index. The virtual node index must fall between 0 and some small maximum. As a result, a node cannot easily control the choice of its own Chord ID.

When a new node n joins the system, some existing nodes may decide to add n to their finger tables. As part of this process, each such existing node sends a message to n 's claimed IP address containing a nonce. If the node at that IP address admits to having n 's ID, and the claimed IP address and virtual node index hash to the ID, then the existing node accepts n .

With this defense in place, an attacker would have to control roughly as many IP addresses as there are total other nodes in the Chord system in order to have a good chance of targeting arbitrary blocks. However, owners of large blocks of IP address space tend to be more easily identifiable (and less likely to be malicious) than individuals.

Chapter 5

DHash Layer

The CFS DHash layer stores and retrieves uniquely identified blocks, and handles distribution, replication, and caching of those blocks. DHash uses Chord to help it locate blocks.

Table 5.1 shows the API that the DHash layer exposes. The CFS file system client layer uses `get` to implement application requests to open files, read files, navigate directories, etc. Publishers of data use a special application that inserts or updates a CFS file system using the `put_h` and `put_s` calls.

5.1 Replication

DHash replicates each block on k CFS servers to increase availability, maintains the k replicas automatically as servers come and go, and places the replicas in a way that clients can easily find them.

DHash places a block's replicas at the k servers immediately after the block's successor on the Chord ring (see Figure 5-1). DHash can easily find the identities of these servers from Chord's r -entry successor list. CFS must be configured so that $r \geq k$. This placement of replicas means that, after a block's successor server fails, the block is immediately available at the block's new successor.

The DHash software in a block's successor server manages replication of that block by making sure that all k of its successor servers have a copy of the block at all times. If the successor server fails, the block's new successor assumes responsibility for the block. The new successor is able to determine that it should assume responsibility by a strictly local examination of routing tables maintained by the Chord layer: if a key is between the node's

Function	Description
<code>put_h(block)</code>	Computes the block's key by hashing its contents, and sends it to the key's successor server for storage.
<code>put_s(block, pubkey)</code>	Stores or updates a signed block; used for root blocks. The block must be signed with the given public key. The block's Chord key will be the hash of <code>pubkey</code> .
<code>get(key)</code>	Fetches and returns the block associated with the specified Chord key.

Table 5.1: DHash client API; exposed to client file system software.




Figure 5-1: The placement of an example block's replicas and cached copies around the Chord identifier ring. The block's ID is shown with a tick mark. The block is stored at the successor of its ID, the server denoted with the square. The block is replicated at the successor's immediate successors (the circles). The hops of a typical lookup path for the block are shown with arrows; the block may be cached at the servers along a possible lookup path (the triangles).

predecessor and its own ID, the node is now responsible for maintaining replicas of the key.

The value of this replication scheme depends in part on the independence of failure and unreachability among a block's k replica servers. Servers close to each other on the ID ring are not likely to be physically close to each other, since a server's ID is based on a hash of its IP address. This provides the desired independence of failure.

CFS could save space by storing coded pieces of blocks rather than whole-block replicas, using an algorithm such as IDA [28]. CFS doesn't use coding, because storage space is not expected to be a highly constrained resource.

5.2 Caching

DHash caches blocks to avoid overloading servers that hold popular data. Each DHash layer sets aside a fixed amount of disk storage for its cache. When a CFS client looks up a block key, it performs a Chord lookup, visiting intermediate CFS servers with IDs successively closer to that of the key's successor (see Figure 5-1). At each step, the client asks the intermediate server whether it has the desired block cached. Eventually the client arrives at either the key's successor or at an intermediate server with a cached copy. The client then sends a copy of the block to the last server it contacted along the lookup path.

Since a Chord lookup takes shorter and shorter hops in ID space as it gets closer to the target, lookups from different clients for the same block will tend to visit the same servers late in the lookup. As a result, the policy of caching blocks along the lookup path is likely to be effective.

By sending a block to the last server contacted we maximize the probability that the cached block will be utilized by future lookups while minimizing the effort spent distributing blocks to caches: the final server is closest to the target and is most likely to be contacted again during a search for the same block. This strategy has the additional property that it spreads cached copies at a distance away from a key's successor in proportion to the popularity of the key.

DHash replaces cached blocks in least-recently-used order. Copies of a block at servers with IDs far from the block's successor are likely to be discarded first, since clients are least likely to stumble upon them. This has the effect of preserving the cached copies close to the successor, and expands and contracts the degree of caching for each block according to its popularity.

While caching and replication are conceptually similar, DHash provides them as distinct mechanisms. DHash stores replicas in predictable places, so it can ensure that enough replicas always exist. In contrast, the number of cached copies cannot easily be counted, and might fall to zero. If fault-tolerance were achieved solely through cached copies, an unpopular block might simply disappear along with its last cached copy.

CFS avoids most cache consistency problems because blocks are keyed by content hashes. Root blocks, however, use public keys as identifiers; a publisher can change a root block by inserting a new one signed with the corresponding private key. This means that cached root blocks may become stale, causing some clients to see an old, but internally consistent, file system. A client can check the freshness of a cached root block [10] to decide whether to look for a newer version. Non-root blocks that no longer have any references to them will eventually be eliminated from caches by LRU replacement.

The caching mechanism described here was implemented to avoid overtaxing servers holding popular documents; the traditional use of caching, however, has been to improve performance. The CFS client (see Section 7.3) caches file blocks in this manner.

5.3 Load Balance

DHash spreads blocks evenly around the ID space, since the content hash function uniformly distributes block IDs. If each CFS server had one ID, the fact that IDs are uniformly distributed would mean that every server would carry roughly the same storage burden. This is not desirable, since different servers may have different storage and network capacities. In addition, even uniform distribution doesn't produce perfect load balance; the maximum storage burden is likely to be about $\log(N)$ times the average due to irregular spacing between server IDs [13].

To accommodate heterogeneous server capacities, CFS uses the notion (from [13]) of a real server acting as multiple *virtual servers*. The CFS protocol operates at the virtual server level. A virtual server uses a Chord ID that is derived from hashing both the real server's IP address and the index of the virtual server within the real server.

A CFS server administrator configures the server with a number of virtual servers in rough proportion to the server's storage and network capacity. This number can be adjusted from time to time to reflect observed load levels.

Use of virtual servers could potentially increase the number of hops in a Chord lookup by increasing the number of nodes participating in the protocol. CFS avoids this expense by allowing virtual servers on the same physical server to examine each others' tables; the fact that these virtual servers can take short-cuts through each others' routing tables exactly compensates for the increased number of servers.

CFS could potentially vary the number of virtual servers per real server adaptively, based on current load. Under high load, a real server could delete some of its virtual servers; under low load, a server could create additional virtual servers. Any such algorithm would need to be designed for stability under high load. If a server is overloaded because the CFS system as a whole is overloaded, then automatically deleting virtual servers might cause a cascade

of such deletions.

5.4 Quotas

The most damaging technical form of abuse that CFS is likely to encounter is malicious injection of large quantities of data. The aim of such an attack might be to use up all the disk space on the CFS servers, leaving none available for legitimate data. Even a non-malicious user could cause the same kind of problem by accident.

Ideally, CFS would impose per-publisher quotas based on reliable identification of publishers, as is done in the PAST system [32]. Reliable identification usually requires some form of centralized administration, such as a certificate authority. As a decentralized approximation, CFS bases quotas on the IP address of the publisher. For example, if each CFS server limits any one IP address to using 0.1% of its storage, then an attacker would have to mount an attack from about 1,000 machines for it to be successful. This mechanism also limits the storage used by each legitimate publisher to just 0.1%, assuming each publisher uses just one IP address.

This limit is not easy to subvert by simple forging of IP addresses, since CFS servers require that publishers respond to a confirmation request that includes a random nonce, as described in Section 4.3. This approach is weaker than one that requires publishers to have unforgeable identities, but requires no centralized administrative mechanisms.

If each CFS server imposes a fixed per-IP-address quota, then the total amount of storage an IP address can consume will grow linearly with the total number of CFS servers. It may prove desirable to enforce a fixed quota on total storage, which would require the quota imposed by each server to decrease in proportion to the total number of servers. An adaptive limit of this form is possible, using the estimate of the total number of servers that the Chord software maintains (see Section 7.1).

5.5 Updates and Deletion

CFS allows updates, but in a way that allows only the publisher of a file system to modify it. A CFS server will accept a request to store a block under either of two conditions. If the block is marked as a content-hash block, the server will accept the block if the supplied key is equal to the SHA-1 hash of the block's content. If the block is marked as a signed block, the block must be signed by a public key whose SHA-1 hash is the block's CFS key.

The low probability of finding two blocks with the same SHA-1 hash prevents an attacker from changing the block associated with a content-hash key, so no explicit protection is required for most of a file system's blocks. The only sensitive block is a file system's root block, which is signed; its safety depends on the publisher avoiding disclosure of the private key.

CFS does not support an explicit delete operation. Publishers must periodically refresh their blocks if they wish CFS to continue to store them. A CFS server may delete blocks that have not been refreshed recently.

One benefit of CFS's implicit deletion is that it automatically recovers from malicious insertions of large quantities of data. Once the attacker stops inserting or refreshing the data, CFS will gradually delete it.

Chapter 6

Server Selection

In building CFS we added the concept of *server selection* to Chord. Server selection reduces lookup latency by allowing lookups to preferentially contact nodes likely to be nearby in the underlying network. It also directs searches to the lowest latency data block replica.

At each step in *find_predecessor(id)* (Figure 4-1), the node doing the lookup (n) can choose the next hop from a set of nodes. Initially this set is the contents of n 's own routing tables; subsequently the set is the list of nodes returned by the *preceding_node_list* RPC to the most recent hop (m). Node m tells n the measured latency from m to each node in the set; m collected these latencies when it acquired its finger table entries. Different choices of next-hop node will take the query different distances around the ID ring, but impose different RPC latencies; the following calculation seeks to pick the best combination.

Chord attempts to predict the latency for the entire lookup given that the next hop is to node n_i for each n_i in the set of potential next hops. This cost, $C(n_i)$, is calculated as follows:

$$\begin{aligned} C(n_i) &= d_i + \bar{d} \times H(n_i) \\ H(n_i) &= \log \left(\frac{n_i - id}{2^{160}} \times N \right) \end{aligned}$$

$H(n_i)$ is an estimate of the number of Chord hops that would remain after contacting n_i . To calculate $H(n_i)$ we begin with N , node n 's estimate of the total number of Chord nodes in the system, based on the density of nodes nearby on the ID ring. Multiplying N by the fraction of the ring between the proposed node and the target gives an estimate of the number of nodes between the proposed node and the target. The base two logarithm of this quantity is an estimate of the number of hops remaining to complete the lookup since we expect the Chord lookup algorithm to eliminate half the remaining nodes between the current position and the target with each hop. Multiplying this quantity by \bar{d} , the average latency of all the RPCs that node n has ever issued, estimates the time it would take to send RPCs for those hops. Adding d_i , the latency to node n_i as reported by node m , produces a complete cost estimate. Chord uses the node with the minimum $C(n_i)$ as the next hop.

Server selection is also performed in the DHash layer when downloading a block. The placement of block replicas makes it easy for a client to select the replica likely to be fastest to download. The result of the Chord lookup for block id is the identity of the server that immediately precedes id . The client asks this predecessor for its successor list, which will include the identities of the servers holding replicas of block id as well as latency

measurements between the predecessor and these servers. The client then fetches the block from the replica with the lowest reported latency.

One benefit of the proposed server selection method is that no extra measurements are necessary to decide which node is closest; the decisions are made based on latencies observed while building finger tables. However, nodes rely on latency measurements taken by other nodes. This works well when low latencies from nodes a to b , and b to c imply that latency is also low from a to c . Measurements of our Internet test-bed suggest that this transitivity condition often holds [37]. An implementation of the lookup algorithm in which nodes issue RPCs on behalf of other nodes could lessen the impact of a lack of transitivity in the network (we call this style of implementation *recursive*).

6.1 Correctness of server selection

Adding server selection makes it more difficult to reason about the performance of Chord. In this section we attempt to provide some basis for an understanding of why the server selection algorithm presented above is helpful.

First we show that the server selection algorithm does not disturb the $O(\log N)$ bound on hops for Chord lookups when all internode latencies are equal.

Observation: In a stable network, if all internode latencies are equal, and nodes are uniformly distributed in the ID space, a lookup operation requires $O(\log N)$ hops.

Pf: If all transmission latencies are equal, the fitness metric ($C(n_i)$) for each potential node can be simplified to $H(n_i)$, the estimate of the number of hops between the proposed node and the target: since the d_i terms for each node are equal they do not affect any comparison. At each step, the algorithm will pick the smallest $H(n_i)$ which is equivalent to picking the n_i with the largest ID since the estimate assumes that nodes are uniformly distributed in ID space. This is equivalent to the behavior of the algorithm as stated without server selection and the $O(\log N)$ bound for that algorithm applies here.

Next we attempt to show how the strictly local, hop by hop, decisions made by the algorithm reduce lookup latency on average. In particular, we wish to demonstrate that when the algorithm does not choose the same node that the unmodified lookup algorithm would, the choice produces a latency savings under certain assumptions.

Observation In a stable network of uniformly distributed (in ID space) nodes, if a search proceeds to a node whose estimated remaining hop count is not the minimum possible choice, the latency of the search will not be increased assuming that the local estimates for average RPC latency and hops remaining are globally accurate.

Pf: If the algorithm chooses a node with a non-minimal hop count there must exist, for some i, j , the comparison:

$$d_i + \bar{d} \times H_i < d_j + \bar{d} \times H_j$$

where $H_i > H_j$. Since $d_j > d_i$ we can write:

$$d_j - d_i > \bar{d} \times (H_i - H_j)$$

$$\text{Latency savings} > \bar{d} \times \text{Added hops}$$

From this result we conclude that the savings we gain by choosing a node that is closer in the network topology but further from the target outweigh, on average, the time spent to traverse the added distance in the ID space.

In practice server selection produces only a prediction for the best node to choose in the next hop. Variable network conditions and intransitivity in the network make it difficult to maintain accurate estimates for internode latencies d_i , the estimate of remaining hops (H_i) is sensitive to fluctuations in the distribution of nodes through the ID space, and \bar{d} is likely to vary based on a node's network connectivity. The result is that server selection is at best a useful heuristic. Worse, server selection may break the assumption that a lookup operation contacts $O(\log N)$ nodes. It is possible that a lookup running under server selection could make a linear number of hops that are "short" in both ID space and in the underlying network topology.

Chapter 7

Implementation

CFS is implemented in 7,000 lines of C++, including the 3,000 line Chord implementation. It consists of a number of separate programs that run at user level. The programs communicate over UDP using a C++ RPC package provided by the SFS toolkit [19]. A busy CFS server may exchange short messages with a large number of other servers, making the overhead of TCP connection setup unattractive. TCP also requires the allocation of resources (a file descriptor, namely) for each connection. In large systems CFS may exhaust these resources by contacting a large number of nodes in a short period of time. The internal structure of each program is based on asynchronous events and callbacks, rather than threads. Each software layer is implemented as a library with a C++ interface. CFS runs on a variety of UNIX systems including Linux, OpenBSD, and FreeBSD. We are currently working on a port of SFS and CFS to the Apple Macintosh OS X operating system.

7.1 Chord Implementation

The Chord library maintains the routing tables described in Section 4. It exports these tables to the DHash layer, which implements its own integrated version of the Chord lookup algorithm. The implementation uses the SHA-1 cryptographic hash function to produce CFS block identifiers from block contents. This means that block and server identifiers are 160 bits wide.

The Chord implementation maintains a running estimate of the total number of Chord servers, for use in the server selection algorithm described in Section 6. Each server computes the fraction of the ID ring that the r nodes in its successor list cover; let that fraction be f . Then the estimated total number of servers in the system is r/f .

7.2 DHash Implementation

DHash is implemented as a library which depends on Chord. Each DHash instance is associated with a Chord virtual server and communicates with that virtual server through a function call interface. DHash instances on different servers communicate with one another via RPC.

DHash has its own implementation of the Chord lookup algorithm, but relies on the Chord layer to maintain the routing tables. Integrating block lookup into DHash increases its efficiency. If DHash instead called the Chord *find_successor* routine, it would be awkward for DHash to check each server along the lookup path for cached copies of the desired block.

```

// RPC handler on server n. Returns block with ID key,
// or the best next server to talk to.
n.lookup_step(key)
  if key ∈ (stored ∪ cached ∪ replicated) then
    return [COMPLETE, key, datakey]
  else if key ∈ (predecessor, myid)
    return NONEXISTENT
  else if key ∈ (myid, first live successor)
    next_hop = first live successor
  else
    // Find highest server < key in my finger table or successor list.
    next_hop = lookup_closest_pred(key)
    succ_list = {s ∈ {fingers ∪ successors} s.t. s > next_hop}
    return [CONTINUE, next_hop, succ_list]

// RPC handler to ask the Chord software to delete
// server id from the finger list and successor list.
n.alert(id)

// Return the block associated with key, or an error.
// Runs on the server that invokes lookup().
lookup(key)
  p.push(n) // A stack to accumulate the path.
  [status, res] = n.lookup_step(key)
  repeat
    if (status = COMPLETE)
      return res.datakey
    else if (status = CONTINUE)
      if (res.next_hop = p.top)
        // p.top knew no server other than itself.
        return NONEXISTENT
      else if (key ∈ (p.top', p.top])
        // p.top should have had the block.
        return NONEXISTENT
      else // explore next hop
        p.push(res.next_hop)
        [status, res] = res.next_hop.lookup_step(key)
    else if (status = RPC_FAILURE)
      // Try again at previous hop.
      failed = p.pop()
      last = p.top()
      last.alert(failed)
      [status, res] = last.lookup_step(key)
  else
    return NONEXISTENT

```

Figure 7-1: The procedure used by DHash to locate a block.

It would also cost an un-needed round trip time, since both Chord and DHash would need to separately contact the block's successor server.

Pseudo-code for the DHash implementation of *lookup(key)* is shown in Figure 7-1; this is DHash's version of the Chord code shown in Figure 4-1. The function *lookup(key)* returns the data associated with *key* or an error if it cannot be found. The function *lookup* operates by repeatedly invoking the remote procedure *n'.lookup_step(key)* which returns one of three possible values. If the called server (*n'*) stores or caches the data associated with *key*, then *n'.lookup_step(key)* returns that data. If *n'* does not store the data, then *n'.lookup_step(key)* returns instead the closest predecessor of *key* (determined by consulting local routing tables on *n'*). Finally, *n'.lookup_step* returns an error if *n'* is the true successor of the key but does not store its associated data.

If *lookup* tries to contact a failed server, the RPC machinery will return `RPC_FAILURE`. The function *lookup* then backtracks to the previously contacted server, tells it about the failed server with *alert*, and asks it for the next-best predecessor. At some point *lookup(key)* will have contacted a pair of servers on either side of *key*. If there have been server failures, the second server of the pair may not be the *key's* original successor. However, that second server will be the first live successor, and will hold a replica for *key*, assuming that not all of its replicas have failed.

Though the pseudo-code does not show it, the virtual servers on any given physical server look at each others' routing tables and block stores. This allows lookups to progress faster around the ring, and increases the chances of encountering a cached block.

As presented in this section, the DHash lookup procedure is *iterative*; that is, the originating server itself queries each server along the lookup path until it reaches a block's successor. We have also implemented the lookup algorithm in a *recursive* manner. In a recursive implementation each node along the lookup path (as opposed to the originating node) issues an RPC to the next node in the search path. The first node encountered which stores the desired data then responds to the originating node with the data. Simplified pseudocode (no failure recovery is shown) for this implementation style is shown in Figure 7-2

7.3 Client Implementation

The CFS client software layers a file system on top of DHash. CFS exports an ordinary UNIX file system interface by acting as a local NFS server using the SFS user level file system toolkit [19]. A CFS file system client (denoted `chordcd` in Figure 7-3 which shows the major components of the CFS client) services NFS requests from applications via the SFS automounter. The file system client communicates with a local block storage daemon (`chordd`) via a UNIX domain socket to fetch file and meta-data blocks from the network. The `chordcd` daemon returns this data to applications via the SFS automounter. The CFS client holds blocks in a local LRU cache to improve performance and reduce network traffic.

File systems are published by another user-level program which reads a directory tree and injects blocks and meta-data into the block store by communicating with a local chord daemon.

The DHash back end is sufficiently flexible to support a number of different client interfaces. For instance, we are currently implementing a client which acts as a web proxy in order to layer the CFS namespace on top of the name space of the world wide web.

```

// local procedure. Begins a recursive lookup
begin_lookup_recursive(key)
  // a nonce allows us to identify which lookup has completed when the lookup returns
  nonce = new_nonce()
  next = lookup_closest_pred(key)
  next.lookup_recursive(key, this.address, nonce)

// RPC handler on server n. Sends an RPC to the next best server
// or to the originator if n holds the requested data
n.lookup_recursive(key, return_address, nonce)
  if key  $\in$  (stored  $\cup$  cached  $\cup$  replicated) then
    return_address.lookup_done(nonce, datakey)
  else if key  $\in$  (predecessor, myid]
    return_address.lookup_done(nonce, NONEXISTENT)
  else if key  $\in$  (myid, first live successor]
    (first live successor).lookup_recursive(nonce, return_address, key)
  else
    // Find highest server < key in my finger table or successor list.
    next_hop = lookup_closest_pred(key)
    next_hop.lookup_recursive(nonce, return_address, key)

// RPC handler on node n. Receives results of recursive lookup.
n.lookup_done(nonce, result)

```

Figure 7-2: A recursive implementation of the DHash lookup algorithm.




Figure 7-3: The design of the CFS client software. Boxes represent processes. Those in gray are not part of CFS. Arrows indicate communication among processes.

7.4 Using CFS

To publish the directory `my-uncopyrighted-music` a user A issues the `chordpblsh` command at the command prompt of his terminal.

```
A@A.com > chordpblsh -k secret_key_file -d my-uncopyrighted-music
chordpblsh: file system exported under IaLCvdNTyE8wX103EAQ2uMqMdYU
```

The `chordpblsh` command will convert the specified directory into a series of data and meta-data blocks and insert them into the DHash system where they are available to other users. The publishing program will sign the root block of the file system using the user's public key (stored, along with the private key, in `secret_key_file`). The hash of this key is encoded in a base 64 ASCII format and becomes the name of the published file system (in this case the name is `IaLCvdNTyE8wX103EAQ2uMqMdYU`). As file system names are not particularly human-readable, we expect that, as in the SFS system, symbolic links will provide authenticated human readable names for documents.

User B could retrieve those files using CFS by accessing the chord namespace embedded in his local namespace on a different node. In the following example the `mpg123` utility is used to play a music file directly from CFS. Alternatively, the file could have been moved to a local disk with the `cp` command.

```
B@B.org > mpg123 /sfs/chord:IaLCvd...uMqMdYU/Bach/T_and_F_in_D_minor.mp3
```


Chapter 8

Experimental Results

In order to demonstrate the practicality of the CFS design, we present two sets of tests. The first explores CFS performance on a modest number of servers distributed over the Internet, and focuses on real-world client-perceived performance. The second involves larger numbers of servers running on a single machine and focuses on scalability and robustness.

Quotas (Section 5.4) were not implemented in the tested software. Cryptographic verification of updates (Section 5.5) was implemented but not enabled. This has no effect on the results presented here.

Unless noted, all tests were run with caching turned off, with no replication, with just one virtual server per physical server, and with server selection turned off. These defaults allow the effects of these features to be individually illustrated. The experiments involve only block-level DHash operations, with no file-system meta-data; the client software driving the experiments fetches a file by fetching a specified list of block identifiers. Every server maintains a successor list with $2 \log_2(N)$ entries, as mentioned in Section 4, to help maintain ring connectivity. While CFS does not automatically adjust the successor list length to match the number of servers, its robustness is not sensitive to the exact value.

8.1 Real Life

We first determine the performance of the Chord protocol; later experiments will exercise DHash. To determine the effectiveness of server selection: 32 small (8 byte) blocks are inserted into a network of 12 nodes scattered over the Internet and then fetched from each server using the iterative lookup algorithm. The machines used in these experiments are part of the RON testbed [2]. A total of three copies of each block are stored. Each lookup took an average of 2.75 hops. As the Figure 8.1 shows, server selection decreases latency appreciably. It is not surprising that server selection does not produce dramatic savings. Because there are only 12 machines in the network we expect that each node will have few unique finger table entries (several nodes had only two unique entries); as a result the number of potential nodes presented to the server selection algorithm is limited.

	no server selection	server selection
Avg. Latency (msec)	384	277

Figure 8-1: Latencies for 8-byte fetches on the RON testbed.




Figure 8-2: Download speeds achieved while fetching a 1MB file with CFS on the Internet testbed, for a range of pre-fetch window sizes. Each point is the average of the times seen by each testbed machine. One curve includes server selection; the other does not.

The remaining tests described in this section used CFS servers running on a testbed of 10 machines scattered over the Internet. Eight of the machines are at sites spread over the United States (three of those are universities connected to the high-speed Abiline or vBNS backbones), and the other two are in Sweden and South Korea.¹

To test CFS we inserted a 1MB file split into 8K blocks. To test download speed, client software on each machine fetched the entire file. The machines fetched the file one at a time. Three RPCs, on average, were required to fetch each block in these experiments. The client software uses pre-fetch to overlap the lookup and fetching of blocks. The client initially issues a window of some number of parallel block fetches; as each fetch completes, the client starts a new one. Five copies (or about $2 \log N$) of each block are stored in the system.

Figure 8-2 shows the average download speeds, for a range of pre-fetch window sizes, with and without server selection. A block fetch without server selection averages nearly 1000 milliseconds; this explains why the download speed is about 10 KBytes/second when fetching one 8KByte block at a time. Increasing the amount of pre-fetch increases the speed; for example, fetching three blocks at a time yields an average speed of about 30 KBytes/second. Large amounts of pre-fetch are counter-productive, since they can congest the client server's network connection. Future versions of CFS will include dynamic flow control to obtain maximum throughput without causing congestion. Server selection increases performance in this scenario, most likely because a large number of data replicas were available (relative to the number of nodes) from which to select. As a result more lookups were directed to the well connected nodes which are able to handle large pre-fetch windows without congesting their links.

Figure 8-3 shows the distribution of speeds seen by the downloads from the different machines, for different pre-fetch windows, with and without server selection. The distributions of speeds are fairly narrow: every download is likely to require a few blocks from every server, so all downloads see a similar mix of per-block times. The best download

¹These results presented here were obtained from testing on a different set of machines than the similar results presented in "Wide Area Cooperative Storage with CFS," SOSP 2001. The machines used here are not, as a group, as well connected as those used to produce results for the SOSP paper. As a result, overall performance is lower (but still competitive with TCP on the same test bed).




Figure 8-3: Cumulative distribution of the download speeds plotted in Figure 8-2, for various pre-fetch windows. Results with and without server selection are marked “w/ s.s.” and “no s.s.” respectively.

speeds were from a machine at New York University with good connections to multiple backbones. The worst download speeds for small pre-fetch windows were from sites outside the United States, which have high latency to most of the servers. The worst speeds for large amounts of pre-fetch were for fetches from cable modem sites in the United States, which have limited link capacity.

Most of the time, server selection improves download speeds by a modest amount. Sometimes it improves them substantially, usually for downloads initiated by well-connected sites. Sometimes server selection makes download speeds worse, usually for downloads initiated by sites outside the United States.

To show that the CFS download speeds are competitive with other file access protocols, files of various sizes were transferred between every pair of the testbed machines using ordinary TCP. The files were transferred one at a time, one whole file per TCP connection. Figure 8-4 shows the cumulative distribution of transfer speeds over the various machine pairs, for 8 KByte, 64 KByte, and 1.1 MByte files. The wide distributions reflect the wide range of propagation delays and link capacities between different pairs of machines. The best speeds are between well-connected sites on the east coast of the United States. The worst speeds for 8 KByte transfers occur when both end-points are outside the United States; the worst for one-megabyte transfers occur when one endpoint is outside the United States and the other is a cable modem, combining high latency with limited link speed.

CFS with a 56 KByte pre-fetch window achieves speeds competitive with TCP, on average. The CFS speeds generally have a distribution much narrower than those of TCP. This means that users are more likely to see repeatably good performance when fetching files with CFS than when fetching files from, for example, FTP servers.




Figure 8-4: Distribution of download speeds achieved by ordinary TCP between each pair hosts on the Internet testbed, for three file sizes.

8.2 Controlled Experiments

The remainder of the results were obtained from a set of CFS servers running on a single machine and using the local loopback network interface to communicate with each other. These servers act just as if they were on different machines. This arrangement is appropriate for controlled evaluation of CFS's scalability and tolerance to failure.

8.2.1 Server Selection

We evaluated the server selection algorithm in an environment with simulated delay: 32 nodes were run on the same physical machine, each one assigned to one of two “continents” based on the node ID's lowest bit. Modifications to the Chord software delayed RPCs between nodes on the same continent by 20ms and by 180ms for nodes on different continents. To test the performance of server selection, 32 keys were inserted into the system and then retrieved from each node. Five copies of each key were present in the system. The numbers presented below represent the average of elapsed times for the 32 nodes to fetch the keys.

	No SS/Iterative	No SS/Recursive	SS/Iterative	SS/Recursive
Avg. Lookup (ms)	451	520	170	228

Times are presented for both a recursive and iterative implementation of the lookup algorithm (see Section 7 for detailed descriptions of these algorithms.) We expect the recursive algorithm to perform better in a network environment that does not exhibit transitivity since latency measurements are used by the nodes that make them rather than by the issuing node as in the iterative case. The recursive lookup implementation presented here has not been as heavily optimized as the iterative version; future implementations of




Figure 8-5: The average number of RPCs that a client must issue to find a block, as function of the total number of servers. The error bars reflect one standard deviation. This experimental data is linear on a log plot, and thus fits the expectation of a logarithmic growth.

recursive lookup should perform better. On average, server selection reduced latency by at least 50 percent in this simulated environment, confirming its value.

8.2.2 Lookup Cost

Looking up a block of data is expected to require $O(\log(N))$ RPCs. The following experiment verifies this expectation. For each of a range of numbers of servers, 10,000 blocks were inserted into the system. Then 10,000 lookups were done, from a single server, for randomly selected blocks. The number of RPCs required for each lookup was recorded. The averages are plotted in Figure 8-5, along with error bars showing one standard deviation.

The results are linear on a log plot, and thus fit the expectation of logarithmic growth. The actual values are about $\frac{1}{2} \log_2(N)$; for example, with 4096 servers, lookups averaged 6.7 RPCs. The number of RPCs required is determined by the number of bits in which the originating server's ID and the desired block's ID differ [35]; this will average about half of the bits, which accounts for the $\frac{1}{2}$.

8.2.3 Load Balance

One of the main goals of CFS is to balance the load over the servers. CFS achieves load balanced storage by breaking file systems up into many blocks and distributing the blocks over the servers. It further balances storage by placing multiple virtual servers per physical server, each virtual server with its own ID. We expect that $O(\log(N))$ virtual servers per physical server will be sufficient to balance the load reasonably well [35].

Figure 8-6 shows typical distributions of ID space among 64 physical servers for 1, 6, and 24 virtual servers per physical server. The crosses represent an actual distribution of




Figure 8-6: Representative cumulative distributions of the fraction of the key space a server might be responsible for. 64 servers are simulated, each with 1, 6, or 24 virtual servers. The data marked Real is derived from the distribution of 10,000 blocks among 64 servers, each with 6 virtual servers.

10,000 blocks over 64 physical servers each with 6 virtual servers. The desired result is that each server's fraction be 0.016. With only one virtual server per server (i.e., without using virtual servers), some servers would store no blocks, and others would store many times the average. With multiple virtual servers per server, the sum of the parts of the ID space that a server's virtual servers are responsible for is more tightly clustered around the average.

The fact that CFS spreads the storage of blocks across servers means that in many cases the burden of serving the blocks will also be evenly spread. For large files this will be true even if some files are more popular than others, since a file's blocks are widely spread. If the popular data consists of only a few blocks, then the servers that happen to be those blocks' successors will experience high load. The next section describes how caching helps balance the serving load for small files.

8.2.4 Caching

CFS caches blocks along the lookup path. As the initiating server contacts successive servers, each checks whether it already has the desired block cached and returns it if so. Once the initiating server has found the block, it sends a copy to the last server contacted before the block was found; this server adds the block to its cache. This scheme is expected to produce high cache hit rates because the lookup paths for the same block from different sources will tend to intersect as they get closer to block's successor server. Figure 8-7 illustrates that lookup paths rapidly intersect. In this experiment each node in a 100 node system looked up the same key and the IDs of nodes visited to resolve the lookup were recorded. A curve in the figure represents the path of a search through ID space: for each hop the ID of the node contacted at that hop is plotted. Only the curves for the 32 nodes which resolved the query in three hops are shown (for clarity's sake). As one can see from




Figure 8-7: Lookup paths in Chord intersect. Each line in figure (a) above represents the path that a lookup took through the ID space to locate a key. The horizontal axis represents the hop number, the vertical axis is the high three bytes of the ID of the node contacted on that hop. As the figure shows, the set of nodes that distinct searches contact diminishes as the searches near the target (corresponding to ID 238). Only lookup paths of length three are shown for clarity (32 such paths are plotted). This experiment was run in a system with 100 nodes. Figure (b) shows the number of nodes visited by any search at each hop for the lookup paths plotted in (a). For instance, the first hop contains 32 nodes since each lookup was initiated from a different node. After one hop lookup paths have converged such that the 32 searches contacted only 10 distinct nodes.

the figure, searches for the same block initiated from a set of nodes distributed across the address space rapidly converge to a few nodes near the target.

Figure 8-8 illustrates how well caching works. A single block is inserted into a 1,000 server system. Then a sequence of randomly chosen servers fetch the block. The graph shows how the number of RPCs required to fetch the block decreases with the number of cumulative fetches, due to the block being cached in more places. Each plotted point is the average of many sequential fetches. A quirk in the implementation prevents the originating server from checking its own cache, which is why no fetches have an RPC count of zero.

As expected, the RPC counts decrease, since more and more servers have the block cached. The RPC counts decrease significantly after just a few lookups. Figure 8-5 shows that lookups without caching in a 1,000-server system require an average of 4 RPCs, while after 10 lookups only 3 hops are required. The net effect is to improve client-perceived performance and to spread the load of serving small files.

To demonstrate how caching distributes the load of serving blocks, we simulated the effect of a single popular file. Figure 8-9 shows the results of this experiment. In this simulation, 100 hosts each ask for a single block. The figure plots the number of RPCs each machine received. This request distribution imposes a large load on the successor node (whose ID corresponds to 92 on the horizontal axis) when caching is not enabled (dotted line). When caching is enabled, the number of RPCs that the successor fields is dramatically reduced. It is also interesting to note that the total number of messages sent when caching is enabled is strictly less than the number sent without caching. When load is considered in terms of message count, caching does not only spread load, but also reduces it. This effect is due to the fact that all nodes which respond with messages returning cached copies of data would have been contacted to forward search requests had caching been disabled.




Figure 8-8: Impact of caching on successive client fetches of the same block. Each point is the average number of RPCs for 10 successive fetches from randomly chosen servers; the error bars indicate one standard deviation. The system has 1,000 servers.




Figure 8-9: Impact of caching on the number of RPCs a server fields. The data is taken from a single run with 100 servers. The horizontal axis indicates the position of the node's ID in a sorted list. The vertical axis indicates the number of RPCs handled by the corresponding node when a single block (stored at node 93) is fetched from every server with caching enabled. Of interest in this figure is the dramatic reduction of the number of RPCs that node 93 must handle




Figure 8-10: Impact of the number of virtual servers per physical server on the total amount of data that the physical server must store.

8.2.5 Storage Space Control

By varying the number of virtual servers on a physical server, a server's owner can control the amount of data that CFS stores on the server, and thus the amount that the server must serve. Figure 8-10 shows how effective this is. The experiment involves seven physical servers, with 1, 2, 4, 8, 16, 32, 64, and 128 virtual servers, respectively. 10,000 blocks are inserted into the system, and the relationship between how many virtual servers a physical server has and how many blocks it must store is plotted. For example, the physical server with 16 virtual server stores 586 blocks; there are a total of 255 virtual servers, so this is close to the expected value $627 = 10000 \times \frac{16}{255}$. Since the relationship of blocks to virtual servers is linear, an administrator can easily adjust a CFS server's storage consumption.

There is little memory overhead to running many virtual servers to achieve fine-grained control over load. Each virtual server requires its own finger table and successor list, as well as accounting structures for the block store and cache; the total memory footprint of these structures in our unoptimized implementation is less than 10KBytes.

8.2.6 Effect of Failure

After a CFS server fails, some time will pass before the remaining servers react to the failure, by correcting their finger tables and successor pointers and by copying blocks to maintain the desired level of replication. Theorems 1 and 2 suggest that CFS will be able to perform lookups correctly and efficiently before this recovery process starts, even in the face of massive failure.

To test this, 1,000 blocks are inserted into a 1,000-server system. Each block has six replicas (including the main copy stored at the direct successor). After the insertions, a fraction of the servers fail without warning. Before Chord starts rebuilding its routing tables, 1,000 fetches of randomly selected blocks are attempted from a single server. Figure 8-11 shows the fraction of lookups that fail, and Figure 8-12 shows the average RPC count of




Figure 8-11: Fraction of block request failures as a function of the fraction of 1,000 CFS servers that fail. Each data point is the average of 5 experiments involving 1,000 block lookups; the error bars indicate the minimum and maximum results.




Figure 8-12: Average lookup RPC count as a function of the fraction of the CFS servers that fail. There are 1,000 servers before the failures. Each data point is the average of 5 experiments; the error bars indicate the minimum and maximum results.




Figure 8-13: Number of RPC timeouts incurred during 1,000 lookups, as a function of the fraction of the CFS servers that fail. There are 1,000 servers before the failures. Each data point is the average of 5 experiments; the error bars indicate the minimum and maximum results.

the lookups.

No lookups fail when fewer than 20% of the servers fail, and very few when less than 35% fail. The reason for this is that server finger tables and successor lists provide many potential paths to carry a query around the Chord ID ring; if the most desirable finger table entry points to a failed server, CFS uses an entry that points less far around the ring. Lookups start to fail when enough servers fail that some blocks lose all six copies. For example, when 50% of the servers fail, the probability of losing all of a block's replicas is $0.5^6 = 0.016$; this is close to the value 0.013 shown in Figure 8-11. All of the lookup failures encountered in this experiment are due to all of a block's replicas failing; CFS was always able to find a copy of a block if one was available.

Figure 8-12 shows that lookups take about one RPC longer as a result of 50% of the servers failing. The RPC counts do not include attempts to contact failed servers. Lookups take longer after failures because some of the finger table entries required for fast lookups point to failed servers. If half of the finger table entries are not valid, then each RPC makes about half as much progress as expected; but one extra RPC fully corrects this.

Figure 8-13 shows the number of attempts to contact failed servers that occur per lookup, averaged over 1,000 block lookups. After the first time a server decides (by a timeout) that it has used a finger table or successor-list entry that points to a failed server, it does not use that server again until it has been stabilized. Given that massive failures have little effect on the availability of data or the number of RPCs per lookup, users are likely to perceive such failures because of RPC timeouts during lookups. However, Figure 8-13 shows that a typical block lookup shortly after a failure can expect less than one timeout on the way to retrieving the desired block.

These experiments demonstrate that a large fraction of CFS servers can fail without significantly affecting data availability or performance.

Chapter 9

Future Work and Conclusions

CFS is just one example of an application that can be built on top of the Chord/DHash system: we hope to explore a variety of systems based on the work presented here. For instance, the most popular peer-to-peer storage systems today offer a keyword search feature. We are currently working on adding this functionality to CFS by distributing an index of content over the nodes. An alternative interface to CFS might make it more usable as well: the file system interface presented here is flexible but requires the user to run an automounting NFS loopback server. A simpler, web-proxy based UI might be preferable for many applications.

We also hope to explore uses for the Chord system other than file storage. Chord acts analogously to DNS: it maps names to servers. This raises the possibility that Chord could provide a decentralized name service with the additional advantage of not requiring the hierarchy imposed by DNS. Providing an authority for a namespace mapped by Chord (i.e. arbitrating conflicts over a given name or set of names) in a decentralized manner is an open problem, however.

A complete defense against malicious nodes has yet to be devised. We believe that a Chord node can verify the data sent to it by consulting additional nodes, but have yet to design an efficient algorithm to do so. The simple defense against flooding attacks described here could also be improved.

Using UDP as an RPC transport over the wide area network saves resources and avoids the expense of setting up connections but sacrifices the flow control properties that TCP provides. Testing CFS made it clear that an effective flow control protocol for CFS traffic is necessary. The many-to-many nature of CFS connectivity will make the design of such a protocol a challenge.

Despite the fact that the work described here is in the early stages of development, this thesis provides a basis for a promising new file distributed file system: CFS is a highly scalable, available and secure read-only file system. It presents stored data to applications through an ordinary file-system interface. Servers store uninterpreted blocks of data with unique identifiers. Clients retrieve blocks from the servers and interpret them as file systems.

CFS stores blocks using the DHash distributed block store. DHash provides an efficient mechanism for storing and retrieving blocks among a decentralized group of servers. DHash replicates a block along consecutive servers in the identifier space and as a result is robust in the face of the failure of a massive number of servers. DHash achieves load balance by caching blocks along probable lookup paths leading to the block's successor.

DHash uses the Chord distributed lookup protocol to map blocks to servers. This

mapping is dynamic and implicit. As a result, there is no directory information to be updated when the underlying network changes (a small number of nodes must update routing tables when the network membership changes, however). This makes CFS both robust and scalable. CFS uses the replication and caching properties of DHash to achieve availability and load balance. Finally, CFS provides simple but effective protection against a single attacker inserting large amounts of data.

A prototype implementation of CFS has been implemented and evaluated on a controlled Internet-wide test-bed. Future operational deployment will likely uncover opportunities for improvement, but the current results indicate that CFS is a viable large-scale peer-to-peer system.

Bibliography

- [1] Akamai Technologies, Inc. <http://www.akamai.com/>, 2001. Cambridge, MA.
- [2] David Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*, October 2001.
- [3] Anawat Chankhunthod, Peter Danzig, Chuck Neerdaels, Michael Schwartz, and Kurt Worrell. A hierarchical Internet object cache. In *Proc. Usenix Technical Conference*, pages 153–163, January 1996.
- [4] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981.
- [5] Ian Clarke. A distributed decentralised information storage and retrieval system. Master’s thesis, University of Edinburgh, 1999.
- [6] Ian Clarke, Oscar Sandberg, Brandon Wiley, and Theodore Hong. Freenet: A distributed anonymous information storage and retrieval system. In *Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, pages 46–66, July 2000.
- [7] cnn.com. Report: Napster users lose that sharing feeling. *cnn.com*. <http://www.cnn.com/2001/TECH/internet/06/28/napster.usage/>.
- [8] Roger Dingledine, Michael Freedman, and David Molnar. The Free Haven project: Distributed anonymous storage service. In *Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, pages 67–95, July 2000.
- [9] L. Fan, P. Cao, J. Almeida, and A. Broder. Summary cache: a scalable wide-area web-cache sharing protocol. Technical Report 1361, Computer Science Department, University of Wisconsin, Madison, February 1998.
- [10] Kevin Fu, M. Frans Kaashoek, and David Mazières. Fast and secure distributed read-only file system. In *Proceedings of the 4th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 181–196, October 2000.
- [11] Syam Gadde, Jeff Chase, and Michael Rabinovich. A taste of crispy squid. In *Workshop on Internet Server Performance*, pages 129–136, June 1998.
- [12] Gnutella website. <http://gnutella.wego.com>.
- [13] D. Karger, E. Lehman, T. Leighton, M. Levine, D. Lewin, and R. Panigrahy. Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the

- world wide web. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 654–663, May 1997.
- [14] Idit Keidar, Jeremy Sussman, Keith Marzullo, and Danny Dolev. Moshe: A group membership service for WANs. Technical Report MIT-LCS-TM-593a, 1999.
 - [15] John Kubiawicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishna Gummadi, Sean Rhea, Hakim Weatherspoon, Westley Weimer, Chris Wells, and Ben Zhao. Oceanstore: An architecture for global-scale persistent storage. In *Proceedings of the Ninth international Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 190–201, November 2000.
 - [16] Daniel Lewin. Consistent hashing and random trees: Algorithms for caching in distributed networks. Master’s thesis, MIT, 1998.
 - [17] Witold Litwin, Marie-Anna Neimat, and Donovan A. Schneider. LH* — a scalable, distributed data structure. *ACM Transactions on Database Systems*, 21(4):480–525, 1996.
 - [18] Radhika Malpani, Jacob Lorch, and David Berger. Making world wide web caching servers cooperate. In *Fourth International World Wide Web Conference*, pages 107–110, 1995.
 - [19] David Mazières. A toolkit for user-level file systems. In *Proc. Usenix Technical Conference*, pages 261–274, June 2001.
 - [20] David Mazières, Michael Kaminsky, M. Frans Kaashoek, and Emmett Witchel. Separating key management from file system security. In *Proceedings of the 17th ACM Symposium on Operating Systems Principles (SOSP)*, pages 124–139, December 1999.
 - [21] Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *Advances in Cryptology—CRYPTO ’87*, volume 293 of *Lecture Notes in Computer Science*, pages 369–378, Berlin, 1987. Springer-Verlag.
 - [22] Mojo nation documentation. <http://www.mojonation.net/docs/>.
 - [23] Morpheus website. <http://www.musiccity.com>.
 - [24] Napster. <http://www.napster.com>.
 - [25] Ohaha. <http://www.ohaha.com/design.html>, as of June 17, 2001, the Ohaha application is no longer available.
 - [26] Andy Oram, editor. *Peer-to-Peer: Harnessing the Power of Disruptive Computation*. O’Reilly & Associates, 2001.
 - [27] C. Plaxton, R. Rajaraman, and A. Richa. Accessing nearby copies of replicated objects in a distributed environment. In *Proceedings of the ACM SPAA*, pages 311–320, June 1997.
 - [28] M. Rabin. Efficient dispersal of information for security, load balancing, and fault tolerance. *Journal of the ACM*, 36(2):335–348, 1989.

- [29] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *Proc. ACM SIGCOMM*, San Diego, 2001.
- [30] Michael Reiter and Aviel Rubin. Crowds: Anonymity for web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, November 1998.
- [31] Antony Rowstron and Peter Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *Proceedings of the 18th IFIP/ACM International Conference on Distributed Systems Platforms (Middleware 2001)*, November 2001.
- [32] Antony Rowstron and Peter Druschel. Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*, October 2001.
- [33] Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble. A measurement study of peer-to-peer file sharing systems. Technical Report UW-CSE-01-06-02, University of Washington, June 2001.
- [34] Alex Sherman, David Karger, Andy Berkheimer, Bill Bogstad, Rizwan Dhanidina, Ken Iwamoto, Brian Kim, Luke Matkins, and Yoav Yerushalmi. Web caching with consistent hashing. *Computer Networks*, 31(11–16):1203–1213, May 1999.
- [35] Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proc. ACM SIGCOMM*, San Diego, 2001.
- [36] Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. Technical Report TR-819, MIT, Cambridge, MA, March 2001.
- [37] Tina Tyan. A case study of server selection. Master’s thesis, MIT, September 2001.
- [38] Marc Waldman, Aviel Rubin, and Lorrie Faith Cranor. Publius: A robust, tamper-evident, censorship-resistant, web publishing system. In *Proc. 9th USENIX Security Symposium*, pages 59–72, August 2000.
- [39] Ben Zhao, John Kubiawicz, and Anthony Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical Report UCB/CSD-01-1141, Computer Science Division, U. C. Berkeley, April 2001.