Improving Storage System Availability With D-GRAID

MUTHIAN SIVATHANU, VIJAYAN PRABHAKARAN, ANDREA C. ARPACI-DUSSEAU, and REMZI H. ARPACI-DUSSEAU
University of Wisconsin, Madison

We present the design, implementation, and evaluation of D-GRAID, a gracefully degrading and quickly recovering RAID storage array. D-GRAID ensures that most files within the file system remain available even when an unexpectedly high number of faults occur. D-GRAID achieves high availability through aggressive replication of semantically critical data, and fault-isolated placement of logically related data. D-GRAID also recovers from failures quickly, restoring only live file system data to a hot spare. Both graceful degradation and live-block recovery are implemented in a prototype SCSI-based storage system underneath unmodified file systems, demonstrating that powerful "file-system like" functionality can be implemented within a "semantically smart" disk system behind a narrow block-based interface.

Categories and Subject Descriptors: D.4.2 [**Operating Systems**]: Storage Management—*Secondary storage*; D.4.5 [**Operating Systems**]: Reliability—*Fault tolerance*

General Terms: Design, Algorithms, Reliability

Additional Key Words and Phrases: Disk array, RAID, Block-based storage, fault isolation, file systems, smart disks

1. INTRODUCTION

> "If a tree falls in the forest and no one hears it, does it make a sound?" *George Berkeley*

Storage systems comprised of multiple disks are the backbone of modern computing centers, and when the storage system is down, the entire center can grind to a halt. Downtime is clearly expensive; for example, in the on-line business world, millions of dollars per hour are lost when systems are not available [Keeton and Wilkes 2002; Patterson 2002].

This work is sponsored by NSF CCR-0092840, CCR-0133456, CCR-0098274, NGS-0103670, ITR-0086044, ITR-0325267, IBM, EMC, and the Wisconsin Alumni Research Foundation.
An earlier version of this article appeared in *Proceedings of the 3rd USENIX Symposium on File and Storage Technologies, 2004* (FAST '04, San Francisco, CA).
Authors' addresses: 7358 Computer Sciences and Statistics, University of Wisconsin, Madison, 1210 Dayton Street, Madison, WI 53706; email: {muthian,vijayan,dusseau,remzi}@cs.wisc.edu.

Storage system *availability* is formally defined as the mean time between failure (MTBF) divided by the sum of the MTBF and the mean time to recovery (MTTR): $\frac{MTBF}{MTBF+MTTR}$ [Gray 1987]. Hence, in order to improve availability, one can either increase the MTBF or decrease the MTTR. Not surprisingly, researchers have studied both of these components of storage availability.

To increase the time between failures of a large storage array, data redundancy techniques can be applied [Bitton and Gray 1988; Burkhard and Menon 1993; Chen et al. 1994; Gray et al. 1990; Hsiao and DeWitt 1990; Orji and Solworth 1993; Park and Balasubramanian 1986; Patterson et al. 1988; Savage and Wilkes 1996; Wilkes et al. 1996]. By keeping multiple copies of blocks, or through more sophisticated redundancy schemes such as parity-encoding, storage systems can tolerate a (small) fixed number of faults. To decrease the time to recovery, "hot spares" can be employed [Holland et al. 1993; Menon and Mattson 1992; Park and Balasubramanian 1986; Reddy and Banerjee 1991]; when a failure occurs, a spare disk is activated and filled with reconstructed data, returning the system to normal operating mode relatively quickly.

1.1 The Problem: Reduced Availability due to Semantic Ignorance

Although various techniques have been proposed to improve storage availability, the narrow interface between file systems and storage [Ganger 2001] has curtailed opportunities for improving MTBF and MTTR. For example, RAID redundancy schemes typically export a simple failure model; if D or fewer disks fail, the RAID continues to operate correctly, but if more than D disks fail, the RAID is entirely unavailable until the problem is corrected, perhaps via a time-consuming restore from tape. In most RAID schemes, D is small (often 1); thus, even when most disks are working, users observe a failed disk system. This "availability cliff" is a result of the storage system laying out blocks oblivious of their semantic importance or relationship; most files become corrupted or inaccessible after just one extra disk failure. Further, because the storage array has no information on which blocks are live in the file system, the recovery process must restore all blocks in the disk. This unnecessary work slows recovery and reduces availability.

An ideal storage array fails gracefully: if $\frac{1}{N}$th of the disks of the system are down, at most $\frac{1}{N}$th of the data is unavailable. An ideal array also recovers intelligently, restoring only live data. In effect, more "important" data is less likely to disappear under failure, and such data is restored earlier during recovery. This strategy for data availability stems from Berkeley's observation about falling trees: if a file isn't available, and no process tries to access it before it is recovered, is there truly a failure?

1.2 A Solution: D-GRAID

To explore these concepts and provide a storage array with more graceful failure semantics, we present the design, implementation, and evaluation of D-GRAID, a RAID system that Degrades Gracefully (and recovers quickly). D-GRAID exploits semantic intelligence [Sivathanu et al. 2003] within the disk array to place file system structures across the disks in a fault-contained

manner, analogous to the fault containment techniques found in the Hive operating system [Chapin et al. 1995] and in some distributed file systems [Ji et al. 2000; Saito et al. 2002]. Thus, when an unexpected "double" failure occurs [Gray 1987], D-GRAID continues operation, serving those files that can still be accessed. D-GRAID also utilizes semantic knowledge during recovery; specifically, only blocks that the file system considers live are restored onto a hot spare. Both aspects of D-GRAID combine to improve the effective availability of the storage array. Note that D-GRAID techniques are complementary to existing redundancy schemes; thus, if a storage administrator configures a D-GRAID array to utilize RAID Level 5, any single disk can fail without data loss, and additional failures lead to a proportional fraction of unavailable data.

In this article, we present a prototype implementation of D-GRAID, which we refer to as *Alexander*. Alexander is an example of a semantically-smart disk system [Sivathanu et al. 2003]. Built underneath a narrow block-based SCSI storage interface, such a disk system understands on-disk file system data structures, including the superblock, allocation bitmaps, inodes, directories, and other important structures; this knowledge is central to implementing graceful degradation and quick recovery. Because of their intricate understanding of file system structures and operations, semantically smart arrays are tailored to particular file systems; Alexander currently functions underneath unmodified Linux ext2 and VFAT file systems.

We make three important contributions to semantic disk technology. First, we deepen the understanding of how to build semantically smart disk systems that operate correctly even with imperfect file system knowledge. Second, we demonstrate that such technology can be applied underneath widely varying file systems. Third, we demonstrate that semantic knowledge allows a RAID system to apply different redundancy techniques based on the type of data, thereby improving availability.

1.3 Key Techniques

There are two key aspects to the Alexander implementation of graceful degradation. The first is *selective meta-data replication*, in which Alexander replicates naming and system meta-data structures of the file system to a high degree while using standard redundancy techniques for data. Thus, with a small amount of overhead, excess failures do not render the entire array unavailable. Instead, the entire directory hierarchy can still be traversed, and only some fraction of files will be missing, proportional to the number of missing disks. The second is a *fault-isolated data placement* strategy. To ensure that semantically meaningful data units are available under failure, Alexander places semantically related blocks (e.g., the blocks of a file) within the storage array's unit of fault-containment (e.g., a disk). By observing the natural failure boundaries found within an array, failures make semantically related groups of blocks unavailable, leaving the rest of the file system intact.

Unfortunately, fault-isolated data placement improves availability at a cost; related blocks are no longer striped across the drives, reducing the natural

benefits of parallelism found within most RAID techniques [Ganger et al. 1993]. To remedy this, Alexander also implements *access-driven diffusion* to improve throughput to frequently-accessed files, by spreading a copy of the blocks of "hot" files across the drives of the system. Alexander monitors access to data to determine which files to replicate in this fashion, and finds space for those replicas either in a preconfigured *performance reserve* or opportunistically in the unused portions of the storage system.

We evaluate the availability improvements possible with D-GRAID through trace analysis and simulation, and find that D-GRAID does an excellent job of masking an arbitrary number of failures from most processes by enabling continued access to "important" data. We then evaluate our prototype, Alexander under microbenchmarks and trace-driven workloads. We find that the construction of D-GRAID is feasible; even with imperfect semantic knowledge, powerful functionality can be implemented within a block-based storage array. We also find that the run-time overheads of D-GRAID are small, but that the storage-level CPU costs compared to a standard array are high. We show that access-driven diffusion is crucial for performance, and that live-block recovery is effective when disks are under-utilized. The combination of replication, data placement, and recovery techniques results in a storage system that improves availability while maintaining a high level of performance.

The rest of this article is structured as follows. In Section 2, we present extended motivation, and in Section 3, we discuss related work. We present the design principles of D-GRAID in Section 4. In Section 5, we present trace analysis and simulations, and discuss semantic knowledge in Section 6. In Section 7, we present our prototype implementation, and evaluate our prototype in Section 8. In Section 9, we present custom policies for different *levels* of D-GRAID. We discuss the resilience of D-GRAID to incorrect information in Section 10, and conclude in Section 11.

2. EXTENDED MOTIVATION

In this section, we first discuss the need for graceful degradation during multiple failures, and then describe why a semantically smart disk system is an appropriate locale to incorporate support for such graceful degradation.

2.1 The Case for Graceful Degradation

The motivation for graceful degradation arises from the fact that users and applications often do not require that the entire contents of a volume be present; rather, what matters to them is whether a particular set of files are available.

One question that arises is whether it is realistic to expect a catastrophic failure scenario within a RAID system. For example, in a RAID-5 system, given the high MTBF's reported by disk manufacturers, one might believe that a second disk failure is highly unlikely to occur before the first failed disk is repaired. However, multiple disk failures do occur, for two primary reasons. First, correlated faults are more common in systems than expected [Gribble 2001]. If the RAID has not been carefully designed in an orthogonal manner, a single controller fault or other component error can render a fair number of

disks unavailable [Chen et al. 1994]; such redundant designs are expensive, and therefore may only be found in higher end storage arrays. Second, Gray [1987] pointed out that system administration is the main source of failure in systems. A large percentage of human failures occur during maintenance, where "the maintenance person typed the wrong command or unplugged the wrong module, thereby introducing a double failure" [Gray 1987, p. 6].

Other evidence also suggests that multiple failures can occur. For example, IBM's ServeRAID array controller product includes directions on how to attempt data recovery when multiple disk failures occur within a RAID-5 storage array [IBM 2001]. Within our own organization, data is stored on file servers under RAID-5. In one of our servers, a single disk failed, but the indicator that should have informed administrators of the problem did not do so. The problem was only discovered when a second disk in the array failed; full restore from backup ran for days. In this scenario, graceful degradation would have enabled access to a large fraction of user data during the long restore.

One might think that the best approach to dealing with multiple failures would be to employ a higher level of redundancy [Alvarez et al. 1997; Burkhard and Menon 1993], thus enabling the storage array to tolerate a greater number of failures without loss of data. However, these techniques are often expensive (e.g., three-way data mirroring) or bandwidth-intensive (e.g., more than 6 I/Os per write in a $P + Q$ redundant store). Graceful degradation is complementary to such techniques. Thus, storage administrators could choose the level of redundancy they believe necessary for common case faults; graceful degradation is enacted when a "worse than expected" fault occurs, mitigating its ill effect.

2.2 The Need for Semantically Smart Storage

The basic design principles of D-GRAID apply equally well to various possible implementation alternatives, each with its own tradeoffs. In this subsection, we motivate our decision to implement D-GRAID within a semantically smart disk system. We first discuss the benefits of such an approach, addressing a few obvious concerns. We then compare the semantic disk approach to other alternatives of implementing D-GRAID.

2.2.1 *Benefits of the Semantic Disk Approach.* Implementing new functionality in a semantically smart disk system has the key benefit of enabling wide-scale deployment underneath an unmodified SCSI interface without any OS modification, thus working smoothly with existing file systems and software base. Although there is some desire to evolve the interface between file systems and storage [Gibson et al. 1998], the reality is that current interfaces will likely survive much longer than anticipated. As Bill Joy once said, "Systems may come and go, but protocols live forever". Similarly to modern processors that innovate beneath unchanged instruction sets, a semantic disk-level implementation is nonintrusive on existing infrastructure, thus making a new technology such as D-GRAID more likely to be adopted.

However, because semantically smart storage systems require more detailed knowledge of the file system that is using them, a few concerns arise on the commercial feasibility of such systems. We consider three main concerns.

The first concern that arises is that placing semantic knowledge within the disk system ties the disk system too intimately to the file system above. For example, if the on-disk structure of the file system changes, the storage system may have to change as well. We believe this issue is not likely to be problematic. On-disk formats evolve slowly, for reasons of backward compatibility. For example, the basic structure of FFS-based file systems has not changed since its introduction in 1984, a period of almost 20 years [McKusick et al. 1984]; the Linux ext2 file system, introduced in roughly 1994, has had the exact same layout for its lifetime; finally, the ext3 journaling file system [Ts'o and Tweedie 2002] is backward compatible with ext2 on-disk layout and the new extensions to the FreeBSD file system [Dowse and Malone 2002] are backward compatible as well. We also have evidence that storage vendors are already willing to maintain and support software specific to a file system; for example, the EMC Symmetrix storage system [EMC Corporation 2002] comes with software that can understand the format of most common file systems.

The second concern is that the storage system needs semantic knowledge for each file system with which it interacts. Fortunately, there are not a large number of file systems that would need to be supported to cover a large fraction of the usage population. If such a semantic storage system is used with a file system that it does not support, the storage system could detect it and turn off its special functionality (e.g., in the case of D-GRAID, revert to normal RAID layout). Such detection can be done by simple techniques such as observing the file system identifier in the partition table.

One final concern that arises is that too much processing will be required within the disk system. We do not believe this to be a major issue, because of the general trend of increasing disk system intelligence [Acharya et al. 1998; Riedel et al. 1998]; as processing power increases, disk systems are likely to contain substantial computational abilities. Indeed, modern storage arrays already exhibit the fruits of Moore's Law; for example, the EMC Symmetrix storage server can be configured with up to 80 processors and 64 GB of RAM [EMC Corporation 2002].

2.2.2 *Comparison with Alternative Approaches.* Although our semantic disk approach has clear benefits as detailed above, it comes with a cost: rediscovering semantic knowledge underneath a modern file system entails a fair amount of complexity.

An alternative approach is to change the interface between file systems and storage, to convey richer information across both layers. For instance, the storage system could expose failure boundaries to the file system [Denehy et al. 2002], and then the file system could explicitly allocate blocks in a fault-isolated manner, placing semantically related blocks together. Alternatively, the file system could tag each write with a logical fault-container ID, which can then be used by the storage system to implement fault-isolated data placement. These techniques, while being conceivably less complex than our approach, have the drawback of being intrusive on existing infrastructure and software base, and requiring wide industry agreement before they can be adopted.

Object-based storage [Gibson et al. 1998] is one such new interface being considered, which makes the file boundaries more visible at the storage layer. However, even with an object-based interface, semantically smart technology will still be relevant to discover semantic relationships across objects, for instance, inferring that a directory object points to a set of file objects which need to be placed within a single fault boundary.

Finally, an approximate version of fault-isolated layout could be implemented in a traditional block based storage system with no semantic understanding. The storage system could simply identify sequences of blocks that are accessed together, and infer that those blocks could be logically related. The main disadvantage with such a black-box approach is that it is quite fragile to concurrent interleavings of independent streams. Further, such a scheme would not be able to identify critical data for purposes of aggressive replication, since such "hot" blocks are most often cached by the file system; thus, frequent reads to them are not visible within the storage system.

3. RELATED WORK

D-GRAID draws on related work from a number of different areas, including distributed file systems and traditional RAID systems. We discuss each in turn.

3.1 Distributed File Systems

Designers of distributed file systems have long ago realized the problems that arise when spreading a directory tree across different machines in a system. For example, Walker et al. discussed the importance of directory namespace replication within the Locus distributed system [Popek et al. 1981]. The Coda mobile file system also takes explicit care with regard to the directory tree [Kistler and Satyanarayanan 1992]. Specifically, if a file is cached, Coda makes sure to cache every directory up to the root of the directory tree. By doing so, Coda can guarantee that a file remains accessible should a disconnection occur. Perhaps an interesting extension to our work would be to reconsider host-based in-memory caching with availability in mind. Also, Slice [Anderson et al. 2002] tries to route namespace operations for all files in a directory to the same server.

More recently, work in wide-area file systems has also reemphasized the importance of the directory tree. For example, the Pangaea file system aggressively replicates the entire tree up to the root on a node when a file is accessed [Saito et al. 2002]. The Island-based file system also points out the need for "fault isolation" but in the context of wide-area storage systems; their "one island principle" is quite similar to fault-isolated placement in D-GRAID [Ji et al. 2000].

Finally, P2P systems such as PAST that place an entire file on a single machine have similar load balancing issues [Rowstron and Druschel 2001]. However, the problem is more difficult in the p2p space due to the constraints of file placement; block migration is much simpler in a centralized storage array.

3.2 Traditional RAID Systems

We also draw on the long history of research in classic RAID systems. From AutoRAID [Wilkes et al. 1996] we learned both that complex functionality could

be embedded within a modern storage array, and that background activity could be utilized successfully in such an environment. From AFRAID [Savage and Wilkes 1996], we learned that there could be a flexible tradeoff between performance and reliability, and the value of delaying updates.

Much of RAID research has focused on different redundancy schemes. While early work stressed the ability to tolerate single-disk failures [Bitton and Gray 1988; Park and Balasubramanian 1986; Patterson et al. 1988], later research introduced the notion of tolerating multiple-disk failures within an array [Alvarez et al. 1997; Burkhard and Menon 1993]. We stress that our work is complementary to this line of research; traditional techniques can be used to ensure full file system availability up to a certain number of failures, and D-GRAID techniques ensure graceful degradation under additional failures. A related approach is parity striping [Gray et al. 1990], which stripes only the parity and not data; while parity striping would achieve a primitive form of fault isolation, the layout is still oblivious of the semantics of the data; blocks will have the same level of redundancy irrespective of their importance (i.e., meta-data vs. data), so multiple failures could still make the entire file system inaccessible. A number of earlier works have also emphasized the importance of hot sparing to speed recovery time in RAID arrays [Holland et al. 1993; Menon and Mattson 1992; Park and Balasubramanian 1986]. Our work on semantic recovery is also complementary to those approaches.

Finally, note that term *graceful degradation* is sometimes used to refer to the performance characteristics of redundant disk systems under failure [Hsiao and DeWitt 1990; Reddy and Banerjee 1991]. This type of graceful degradation is different from what we discuss in this article; indeed, none of those systems continues operation when an unexpected number of failures occurs.

4. DESIGN: D-GRAID EXPECTATIONS

We now discuss the design of D-GRAID. We present background information on file systems, the data layout strategy required to enable graceful degradation, the important design issues that arise due to the new layout, and the process of fast recovery.

4.1 File System Background

Semantic knowledge is system specific; therefore, we discuss D-GRAID design and implementation for two widely differing file systems: the Linux ext2 [Ts'o and Tweedie 2002] and Microsoft VFAT [Microsoft Corporation 2000] file systems. Inclusion of VFAT represents a significant contribution compared to previous research, which operated solely underneath UNIX file systems.

The ext2 file system is an intellectual descendant of the Berkeley Fast File System (FFS) [McKusick et al. 1984]. The disk is split into a set of *block groups*, akin to cylinder groups in FFS, each of which contains bitmaps to track inode and data block allocation, inode blocks, and data blocks. Most information about a file, including size and block pointers, are found in the file's inode.

The VFAT file system descends from the world of PC operating systems. In this article, we consider the Linux VFAT implementation of FAT-32. VFAT

operations are centered around the eponymous file allocation table, which contains an entry for each allocatable block in the file system. These entries are used to locate the blocks of a file, in a linked-list fashion, for example, if a file's first block is at address *b*, one can look in entry *b* of the FAT to find the next block of the file, and so forth. An entry can also hold an end-of-file marker or a setting that indicates the block is free. Unlike UNIX file systems, where most information about a file is found in its inode, a VFAT file system spreads this information across the FAT itself and the directory entries; the FAT is used to track which blocks belong to the file, whereas the directory entry contains information like size, permission, and type information.

4.2 Graceful Degradation

To ensure partial availability of data under multiple failures in a RAID array, D-GRAID employs two main techniques. The first is a *fault-isolated data placement* strategy, in which D-GRAID places each "semantically related set of blocks" within a "unit of fault containment" found within the storage array. For simplicity of discussion, we assume that a file is a semantically related set of blocks, and that a single disk is the unit of fault containment. We will generalize the former below, and the latter is easily generalized if there are other failure boundaries that should be observed (e.g., SCSI chains). We refer to the physical disk to which a file belongs as the *home site* for the file. When a particular disk fails, fault-isolated data placement ensures that only files that have that disk as their home site become unavailable, while other files remain accessible as whole files.

The second technique is *selective meta-data replication*, in which D-GRAID replicates naming and system meta-data structures of the file system to a high degree, for example, directory inodes and directory data in a UNIX file system. D-GRAID thus ensures that all live data is reachable and not orphaned due to multiple failures. The entire directory hierarchy remains traversable, and the fraction of missing user data is proportional to the number of failed disks.

Thus, D-GRAID lays out logical file system blocks in such a way that the availability of a single file depends on as few disks as possible. In a traditional RAID array, this dependence set is normally the entire set of disks in the group, thereby leading to entire file system unavailability under an unexpected failure. A UNIX-centric example of typical layout, fault-isolated data placement, and selective meta-data replication is depicted in Figure 1. Note that for the techniques in D-GRAID to work, a meaningful subset of the file system must be laid out within a single D-GRAID array. For example, if the file system is striped across multiple D-GRAID arrays, no single array will have a meaningful view of the file system. In such a scenario, D-GRAID can be run at the logical volume manager level, viewing each of the arrays as a single disk; the same techniques remain relevant.

Because D-GRAID treats each file system block type differently, the traditional RAID taxonomy is no longer adequate in describing how D-GRAID behaves. Instead, a more fine-grained notion of a RAID level is required, as D-GRAID may employ different redundancy techniques for different types of

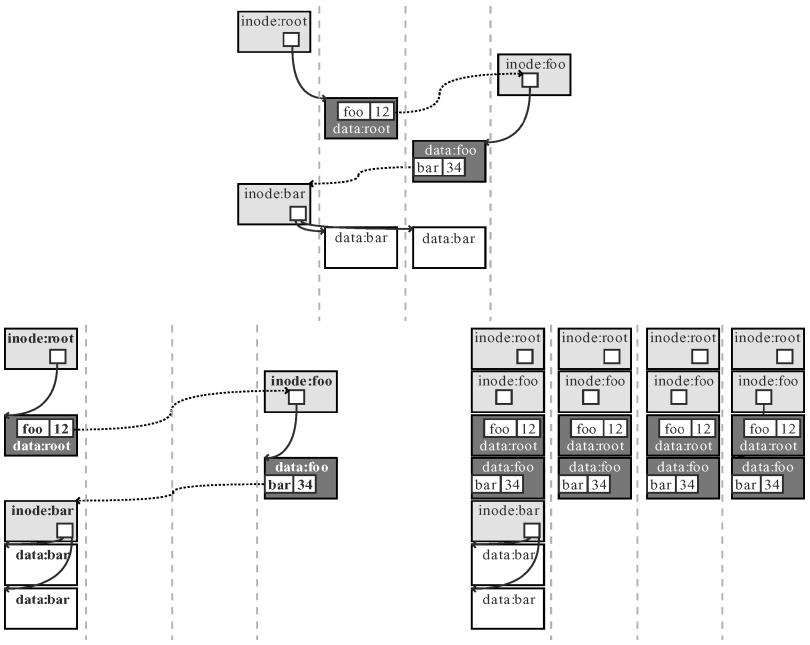

Fig. 1. A comparison of layout schemes. These different parts of this figure depict different layouts of a file "/foo/bar" in a UNIX file system starting at the root inode and following down the directory tree to the file data. Each vertical column represents a disk. For simplicity, the example assumes no data redundancy for user file data. On the top is a typical file system layout on a non-D-GRAID disk system; because blocks (and therefore pointers) are spread throughout the file system, any single fault will render the blocks of the file "bar" inaccessible. The left figure in the bottom is a fault-isolated data placement of files and directories. In this scenario, if one can access the inode of a file, one can access its data (indirect pointer blocks would also be constrained within the same disk). Finally, in the bottom right is an example of selective meta-data replication. By replicating directory inodes and directory blocks, D-GRAID can guarantee that users can get to all files that are available. Some of the requisite pointers have been removed from the rightmost figure for simplicity. Color codes are white for user data, light shaded for inodes, and dark shaded for directory data.

data. For example, D-GRAID commonly employs n-way mirroring for naming and system meta-data, whereas it uses standard redundancy techniques, such as mirroring or parity encoding (e.g., RAID-5), for user data. Note that n, a value under administrative control, determines the number of failures under which D-GRAID will degrade gracefully. In Section 5, we will explore how data availability degrades under varying levels of namespace replication.

4.3 Design Considerations

The layout and replication techniques required to enable graceful degradation introduce a number of design issues. We highlight the major challenges that arise.

4.3.1 *Semantically Related Blocks.* With fault-isolated data placement, D-GRAID places a logical unit of file system data (e.g., a file) within a fault-isolated container (e.g., a disk). Which blocks D-GRAID considers "related" thus

determines which data remains available under failure. The most basic approach is *file-based* grouping, in which a single file (including its data blocks, inode, and indirect pointers) is treated as the logical unit of data; however, with this technique a user may find that some files in a directory are unavailable while others are not, which may cause frustration and confusion. Other groupings preserve more meaningful portions of the file system volume under failure. With *directory-based* grouping, D-GRAID ensures that the files of a directory are all placed within the same unit of fault containment. Less automated options are also possible, allowing users to specify arbitrary semantic groupings which D-GRAID then treats as a unit.

4.3.2 *Load Balance.* With fault-isolated placement, instead of placing blocks of a file across many disks, the blocks are isolated within a single home site. Isolated placement improves availability but introduces the problem of load balancing, which has both space and time components.

In terms of space, the total utilized space in each disk should be maintained at roughly the same level, so that, when a fraction of disks fail, roughly the same fraction of data becomes unavailable. Such balancing can be addressed in the foreground (i.e., when data is first allocated), the background (i.e., with migration), or both. Files (or directories) larger than the amount of free space in a single disk can be handled either with a potentially expensive reorganization or by reserving large extents of free space on a subset of drives. Files that are larger than a single disk must be split across disks.

More pressing are the performance problems introduced by fault-isolated data placement. Previous work has indicated that striping of data across disks is better for performance even compared to sophisticated file placement algorithms [Ganger et al. 1993; Wolf 1989]. Thus, D-GRAID makes additional copies of user data that are spread across the drives of the system, a process which we call *access-driven diffusion*. Whereas standard D-GRAID data placement is optimized for availability, access-driven diffusion increases performance for those files that are frequently accessed. Not surprisingly, access-driven diffusion introduces policy decisions into D-GRAID, including where to place replicas that are made for performance, which files to replicate, and when to create the replicas.

4.3.3 *Meta-Data Replication Level.* The degree of meta-data replication within D-GRAID determines how resilient it is to excessive failures. Thus, a high degree of replication is desirable. Unfortunately, meta-data replication comes with costs, both in terms of space and time. For space overheads, the tradeoffs are obvious: more replicas imply more resiliency. One difference between traditional RAID and D-GRAID is that the amount of space needed for replication of naming and system meta-data is dependent on usage, that is, a volume with more directories induces a greater amount of overhead. For time overheads, a higher degree of replication implies lowered write performance for naming and system meta-data operations. However, others have observed that there is a lack of update activity at higher levels in the directory tree [Popek et al. 1981], and lazy update propagation can be employed to reduce costs [Savage and Wilkes 1996].

4.4 Fast Recovery

Because the main design goal of D-GRAID is to ensure higher availability, fast recovery from failure is also critical. The most straightforward optimization available with D-GRAID is to recover only "live" file system data. Assume we are restoring data from a live mirror onto a hot spare; in the straightforward approach, D-GRAID simply scans the source disk for live blocks, examining appropriate file system structures to determine which blocks to restore. This process is readily generalized to more complex redundancy encodings. D-GRAID can potentially prioritize recovery in a number of ways, for example, by restoring certain "important" files first, where importance could be domain specific (e.g., files in `/etc`) or indicated by users in a manner similar to the hoarding database in Coda [Kistler and Satyanarayanan 1992].

5. EXPLORING GRACEFUL DEGRADATION

In this section, we use simulation and trace analysis to evaluate the potential effectiveness of graceful degradation and the impact of different semantic grouping techniques. We first quantify the space overheads of D-GRAID. Then we demonstrate the ability of D-GRAID to provide continued access to a proportional fraction of meaningful data after arbitrary number of failures. More importantly, we then demonstrate how D-GRAID can hide failures from users by replicating "important" data. The simulations use file system traces collected from HP Labs [Riedel et al. 2002], and cover 10 days of activity; there are 250 GB of data spread across 18 logical volumes.

5.1 Space Overheads

We first examine the space overheads due to selective meta-data replication that are typical with D-GRAID-style redundancy. We calculate the cost of selective meta-data replication as a percentage overhead, measured across all volumes of the HP trace data when laid out in either the ext2 or the VFAT file system. When running underneath ext2, selective meta-data replication is applied to the superblock, inode and data block bitmaps, and the inode and data blocks of directory files. The blocks replicated in the case of VFAT are those that comprise the FAT and the directory entries. We calculate the highest possible percentage of selective meta-data replication overhead by assuming no replication of user data; if user data is mirrored, the overheads are cut in half.

Table I shows that selective meta-data replication induces only a mild space overhead even under high levels of meta-data redundancy for both the Linux ext2 and VFAT file systems. Even with 16-way redundancy of meta-data, only a space overhead of 8% is incurred in the worst case (VFAT with 1-kB blocks). With increasing block size, while ext2 uses more space (due to internal fragmentation with larger directory blocks), the overheads actually decrease with VFAT. This phenomenon is due to the structure of VFAT; for a fixed-sized file system, as block size grows, the file allocation table itself shrinks, although the blocks that contain directory data grow.

Table I. Space Overhead of Selective Meta-Data Replication (The table shows the space overheads of selective meta-data replication as a percentage of total user data, as the level of naming and system meta-data replication increases. In the leftmost column, the percentage space overhead without any meta-data replication is shown. The next two columns depict the costs of modest (four-way) and paranoid (16-way) schemes. Each row shows the overhead for a particular file system, either ext2 or VFAT, with block size set to 1 kB or 4 kB)

	Level of Replication		
	1-way	4-way	16-way
$ext2_{1KB}$	0.15%	0.60%	2.41%
$ext2_{4KB}$	0.43%	1.71%	6.84%
$VFAT_{1KB}$	0.52%	2.07%	8.29%
$VFAT_{4KB}$	0.50%	2.01%	8.03%

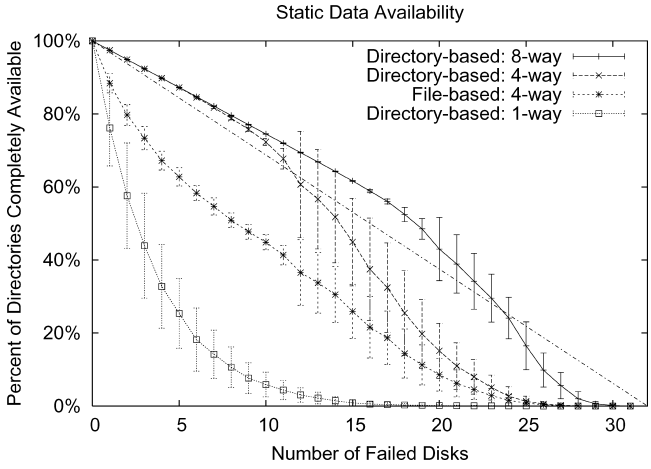

Fig. 2. Static data availability. The percent of entire directories available is shown under increasing disk failures. The simulated system consists of 32 disks, and is loaded with the 250 GB from the HP trace. Two different strategies for semantic grouping are shown: file-based and directory-based. Each line varies the level of replication of namespace meta-data. Each point shows average and deviation across 30 trials, where each trial randomly varies which disks fail.

5.2 Static Availability

We next examine how D-GRAID availability degrades under failure with two different semantic grouping strategies. The first strategy is file-based grouping, which keeps the information associated with a single file within a failure boundary (i.e., a disk); the second is directory-based grouping, which allocates files of a directory together. For this analysis, we place the entire 250 GB of files and directories from the HP trace onto a simulated 32-disk system, remove simulated disks, and measure the percentage of whole directories that are available. We assume no user data redundancy (i.e., D-GRAID Level 0).

Figure 2 shows the percent of directories available, where a directory is available if all of its files are accessible (although subdirectories and their files may not be). From the figure, we observe that graceful degradation works quite well, with the amount of available data proportional to the number of working

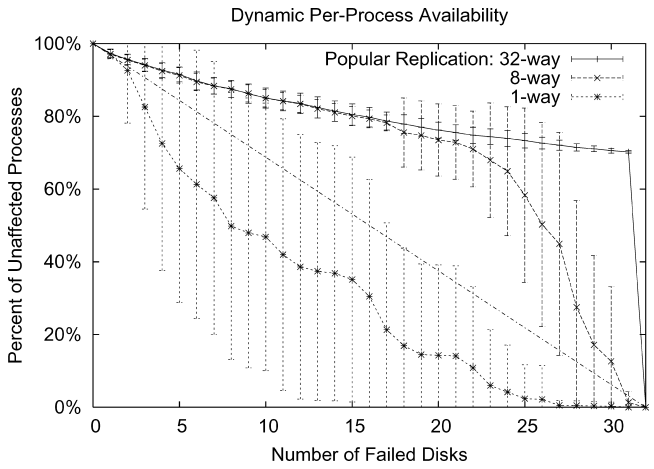

Fig. 3. Dynamic data availability. The figure plots the percent of processes that run unaffected under disk failure from one busy hour from the HP trace. The degree of namespace replication is set aggressively to 32. Each line varies the amount of replication for "popular" directories; one-way implies that those directories are not replicated, whereas eight-way and 32-way show what happens with a modest and extreme amount of replication. Means and deviations of 30 trials are shown.

disks, in contrast to a traditional RAID where a few disk crashes would lead to complete data unavailability. In fact, availability sometimes degrades slightly less than expected from a strict linear fall-off; this is due to a slight imbalance in data placement across disks and within directories. Further, even a modest level of namespace replication (e.g., four-way) leads to very good data availability under failure. We also conclude that with file-based grouping, some files in a directory are likely to "disappear" under failure, leading to user dissatisfaction.

5.3 Dynamic Availability

Finally, by simulating dynamic availability, we examine how often users or applications will be oblivious that D-GRAID is operating in degraded mode. Specifically, we run a portion of the HP trace through a simulator with some number of failed disks, and record what percent of processes observed no I/O failure during the run. Through this experiment, we find that namespace replication is not enough; certain files, which are needed by most processes, must be replicated as well.

In this experiment, we set the degree of namespace replication to 32 (full replication), and vary the level of replication of the contents of popular directories, that is, `/usr/bin`, `/bin`, `/lib`, and a few others. Figure 3 shows that, without replicating the contents of those directories, the percentage of processes that run without ill-effect is lower than expected from our results in Figure 2. However, when those few directories are replicated, the percentage of processes that run to completion under disk failure is much better than expected. The reason for this is clear: a substantial number of processes (e.g., `who`, `ps`, etc.) only require that their executable and a few other libraries be available to run correctly. With popular directory replication, excellent availability under

failure is possible. Fortunately, almost all of the popular files are in "read-only" directories; thus, wide-scale replication will not raise write performance or consistency issues. Also, the space overhead due to popular directory replication is minimal for a reasonably sized file system; for this trace, such directories account for about 143 MB, less than 0.1% of the total file system size.

6. SEMANTIC KNOWLEDGE

We now move toward the construction of a D-GRAID prototype underneath a block-based SCSI-like interface. The enabling technology underlying D-GRAID is semantic knowledge [Sivathanu et al. 2003]. Understanding how the file system above utilizes the disk enables D-GRAID to implement both graceful degradation under failure and quick recovery. The exact details of acquiring semantic knowledge within a disk or RAID system have been described elsewhere [Sivathanu et al. 2003]; here we just assume that a basic understanding of file system layout and structures is available within the storage system. Specifically, we assume that D-GRAID has static knowledge of file system layout, including which regions on disk are used for which block types and the contents of specific block types, for example, the fields of an inode. As we describe below, D-GRAID builds upon this basic knowledge to infer more detailed dynamic information about the file system.

6.1 File System Behaviors

In this article, we extend understanding of semantically smart disks by presenting techniques to handle more general file system behaviors. Previous work required the file system to be mounted synchronously when implementing complex functionality within the disk; we relax that requirement. We now describe the set of typical file system properties that are important from the viewpoint of a semantically smart disk. We believe that many, if not all, modern file systems adhere to these behaviors.

First, blocks in a file system can be dynamically typed, that is, the file system can locate different types of blocks at the same physical location on disk over the lifetime of the file system. For example, in a UNIX file system, a block in the data region can be a user-data block, an indirect-pointer block, or a directory-data block. Second, a file system can delay updates to disk; delayed writes at the file system facilitate batching of small writes in memory and suppressing of writes to files that are subsequently deleted. Third, as a consequence of delayed writes, the order in which the file system actually writes data to disk can be arbitrary. Although certain file systems order writes carefully [Ganger et al. 2000], to remain general, we do not make any such assumptions on the ordering. Note that these properties are identified for practical reasons: the Linux ext2 file system exhibits all the aforementioned behaviors.

6.2 Accuracy of Information

Our assumptions about general file system behavior imply that the storage system cannot accurately classify the type of each block. Block classification is straightforward when the type of the block depends upon its location on disk.

For example, in the Berkeley Fast File System (FFS) [McKusick et al. 1984], the regions of disk that store inodes are fixed at file system creation; thus, any traffic to those regions is known to contain inodes.

However, type information is sometimes spread across multiple blocks. For example, a block filled with indirect pointers can only be identified as such by observing the corresponding inode, specifically that the inode's indirect pointer field contains the address of the given indirect block. More formally, to identify an indirect block B, the semantic disk must look for the inode that has block B in its indirect pointer field. Thus, when the relevant inode block I_B is written to disk, the disk infers that B is an indirect block, and when it later observes block B written, it uses this information to classify and treat the block as an indirect block. However, due to the delayed write and reordering behavior of the file system, it is possible that in the time between the disk writes of I_B and B, block B was freed from the original inode and was reallocated to another inode with a different type, that is, as a normal data block. The disk does not know this since the operations took place in memory and were not reflected to disk. Thus, the inference made by the semantic disk on the block type could be wrong due to the inherent staleness of the information tracked. Implementing a correct system despite potentially inaccurate inferences is one of the challenges we address in this article.

7. IMPLEMENTATION: MAKING D-GRAID

We now discuss the prototype implementation of D-GRAID known as *Alexander*. Alexander uses fault-isolated data placement and selective meta-data replication to provide graceful degradation under failure, and employs access-driven diffusion to correct the performance problems introduced by availability-oriented layout. Currently, Alexander replicates namespace and system meta-data to an administrator-controlled value (for example, 4 or 8), and stores user data in either a RAID-0 or RAID-1 manner; we refer to those systems as D-GRAID *Levels 0* and *1*, respectively. We are currently pursuing a D-GRAID Level 5 implementation, which uses log-structuring [Rosenblum and Ousterhout 1992] to avoid the small-write problem that is exacerbated by fault-isolated data placement.

In this section, we present the implementation of graceful degradation and live-block recovery, with most of the complexity (and hence discussion) centered around graceful degradation. For simplicity of exposition, we focus on the construction of Alexander underneath the Linux ext2 file system. At the end of the section, we discuss differences in our implementation underneath VFAT.

7.1 Graceful Degradation

We now present an overview of the basic operation of graceful degradation within Alexander. We first describe the simple cases before proceeding to the more intricate aspects of the implementation.

7.1.1 *The Indirection Map.* Similarly to any other SCSI-based RAID system, Alexander presents host systems with a linear logical block address space.

Internally, Alexander must place blocks so as to facilitate graceful degradation. Thus, to control placement, Alexander introduces a transparent level of indirection between the logical array used by the file system and physical placement onto the disks via the *indirection map (imap)*; similar structures have been used by others [English and Stepanov 1992; Wang et al. 1999; Wilkes et al. 1996]. Unlike most of these other systems, this imap only maps every *live* logical file system block to its replica list, that is, all its physical locations. All *unmapped* blocks are considered free and are candidates for use by D-GRAID.

7.1.2 *Reads.* Handling block read requests at the D-GRAID level is straightforward. Given the logical address of the block, Alexander looks in the imap to find the replica list and issues the read request to one of its replicas. The choice of which replica to read from can be based on various criteria [Wilkes et al. 1996]; currently Alexander uses a randomized selection. However, in the presence of access-driven diffusion, the diffused copy is always given preference over the fault-isolated copy.

7.1.3 *Writes.* In contrast to reads, write requests are more complex to handle. Exactly how Alexander handles the write request depends on the *type* of the block that is written. Figure 4 depicts some common cases.

If the block is a static meta-data block (e.g., an inode or a bitmap block) that is as yet unmapped, Alexander allocates a physical block in each of the disks where a replica should reside, and writes to all of the copies. Note that Alexander can easily detect static block types such as inode and bitmap blocks underneath many UNIX file systems simply by observing the logical block address.

When an inode block is written, D-GRAID scans the block for newly added inodes; to understand which inodes are new, D-GRAID compares the newly written block with its old copy, a process referred to as *block differencing*. For every new inode, D-GRAID selects a home site to lay out blocks belonging to the inode, and records it in the *inode-to-home-site* hashtable. This selection of home site is done to balance space allocation across physical disks. Currently, D-GRAID uses a greedy approach; it selects the home site with the most free space.

If the write is to an unmapped block in the data region (i.e., a data block, an indirect block, or a directory block), the allocation cannot be done until D-GRAID knows which file the block belongs to, and thus, its actual home site. In such a case, D-GRAID places the block in a *deferred block list* and does not write it to disk until it learns which file the block is associated with. Since a crash before the inode write would make the block inaccessible by the file system anyway, the in-memory deferred block list is not a reliability concern.

D-GRAID also looks for newly added block pointers when an inode (or indirect) block is written. If the newly added block pointer refers to an unmapped block, D-GRAID adds a new entry in the imap, mapping the logical block to a physical block in the home site assigned to the corresponding inode. If any newly added pointer refers to a block in the deferred list, D-GRAID removes the block from the deferred list and issues the write to the appropriate physical block(s). Thus, writes are deferred only for blocks that are written *before* the

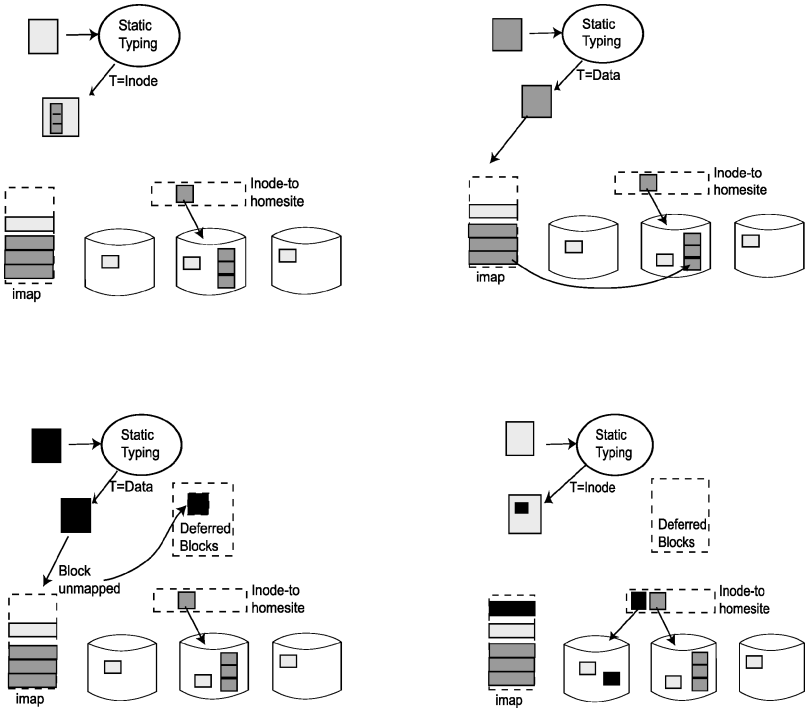

Fig. 4. Anatomy of a `write`. The parts of this figure depicts the control flow during a sequence of `write` operations in Alexander. In the first figure, an inode block is written; Alexander observes the contents of the inode block and identifies the newly added inode. It then selects a home site for the inode and creates physical mappings for the blocks of the inode, in that home site. Also, the inode block is aggressively replicated. In the next part, Alexander observes a write to a data block from the same inode; since it is already mapped, the `write` goes directly to the physical block. In the third part Alexander gets a `write` to an unmapped data block; it therefore defers writing the block, and when Alexander finally observes the corresponding inode (in the fourth part), it creates the relevant mappings, observes that one of its blocks is deferred, and therefore issues the deferred `write` to the relevant home site.

corresponding owner inode blocks. If the inode is written first, subsequent data writes will be already mapped and sent to disk directly.

Another block type of interest that D-GRAID looks for is the data bitmap block. Whenever a data bitmap block is written, D-GRAID scans through it looking for newly freed data blocks. For every such freed block, D-GRAID removes the logical-to-physical mapping if one exists and frees the corresponding physical blocks. Further, if a block that is currently in the deferred list is freed, the block is removed from the deferred list and the write is suppressed; thus, data blocks that are written by the file system but deleted before their corresponding inode is written to disk do not generate extra disk traffic, similarly to optimizations found in many file systems [Rosenblum and Ousterhout 1992]. Removing such blocks from the deferred list is important because, in the case of freed blocks, Alexander may never observe an owning inode. Thus, every deferred block stays in the deferred list for a bounded amount of time, until either an inode owning the block is written, or a bitmap

block indicating deletion of the block is written. The exact duration depends on the delayed write interval of the file system.

7.1.4 *Block Reuse.* We now discuss a few of the more intricate issues involved with implementing graceful degradation. The first such issue is block reuse. As existing files are deleted or truncated and new files are created, blocks that were once part of one file may be reallocated to some other file. Since D-GRAID needs to place blocks onto the correct home site, this reuse of blocks needs to be detected and acted upon. D-GRAID handles block reuse in the following manner: whenever an inode block or an indirect block is written, D-GRAID examines each valid block pointer to see if its physical block mapping matches the home site allocated for the corresponding inode. If not, D-GRAID changes the mapping for the block to the correct home site. However, it is possible that a write to this block (which was made in the context of the new file) went to the old home site, and hence needs to be copied from its old physical location to the new location. Blocks that must be copied are added to a *pending copies list*; a background thread copies the blocks to their new homesite and frees the old physical locations when the copy completes.

7.1.5 *Dealing with Imperfection.* Another difficulty that arises in semantically smart disks underneath typical file systems is that exact knowledge of the type of a dynamically typed block is impossible to obtain, as discussed in Section 6. Thus, Alexander must handle incorrect type classification for data blocks (i.e., file data, directory, and indirect blocks).

For example, D-GRAID must understand the contents of indirect blocks, because it uses the pointers therein to place a file's blocks onto its home site. However, due to lack of perfect knowledge, the fault-isolated placement of a file might be compromised (note that data loss or corruption is not an issue). Our goal in dealing with imperfection is thus to conservatively avoid it when possible, and eventually detect and handle it in all other cases.

Specifically, whenever a block construed to be an indirect block is written, we assume it is a valid indirect block. Thus, for every live pointer in the block, D-GRAID must take some action. There are two cases to consider. In the first case, a pointer could refer to an unmapped logical block. As mentioned before, D-GRAID then creates a new mapping in the home site corresponding to the inode to which the indirect block belongs. If this indirect block (and pointer) is valid, this mapping is the correct mapping. If this indirect block is misclassified (and consequently, the pointer invalid), D-GRAID detects that the block is free when it observes the data bitmap write, at which point the mapping is removed. If the block is allocated to a file before the bitmap is written, D-GRAID detects the reallocation during the inode write corresponding to the new file, creates a new mapping, and copies the data contents to the new home site (as discussed above).

In the second case, a potentially corrupt block pointer could point to an already mapped logical block. As discussed above, this type of block reuse results in a new mapping and copy of the block contents to the new home site. If this indirect block (and hence, the pointer) is valid, this new mapping is the

correct one for the block. If instead the indirect block is a misclassification, Alexander wrongly copies over the data to the new home site. Note that the data is still accessible; however, the original file to which the block belongs, now has one of its blocks in the incorrect home site. Fortunately, this situation is transient, because once the inode of the file is written, D-GRAID detects this as a reallocation and creates a new mapping back to the original home site, thereby restoring its correct mapping. Files which are never accessed again are properly laid out by an infrequent sweep of inodes that looks for rare cases of improper layout.

Thus, without any optimizations, D-GRAID will eventually move data to the correct home site, thus preserving graceful degradation. However, to reduce the number of times such a misclassification occurs, Alexander makes an assumption about the contents of indirect blocks, specifically that they contain some number of valid unique pointers, or null pointers. Alexander can leverage this assumption to greatly reduce the number of misclassifications, by performing an integrity check on each supposed indirect block. The integrity check, which is reminiscent of work on conservative garbage collection [Boehm and Weiser 1988], returns true if all the "pointers" (4-byte words in the block) point to valid data addresses within the volume and all nonnull pointers are unique. Clearly, the set of blocks that pass this integrity check could still be corrupt if the data contents happened to exactly evade our conditions. However, a test run across the data blocks of our local file system indicates that only a small fraction of data blocks (less than 0.1%) would pass the test; only those blocks that pass the test *and* are reallocated from a file data block to an indirect block would be misclassified.[1]

7.1.6 *Access-Driven Diffusion.* Another issue that D-GRAID must address is performance. Fault-isolated data placement improves availability but at the cost of performance. Data accesses to blocks of a large file, or, with directory-based grouping, to files within the same directory, are no longer parallelized. To improve performance, Alexander performs access-driven diffusion, monitoring block accesses to determine which block ranges are "hot," and then "diffusing" those blocks via replication across the disks of the system to enhance parallelism.

Access-driven diffusion can be achieved at both the logical and physical levels of a disk volume. In the logical approach, access to individual files is monitored, and those considered hot are diffused. However, per-file replication fails to capture sequentiality across multiple small files, for example, those in a single directory. Therefore we instead pursue a physical approach, in which Alexander replicates segments of the logical address space across the disks of the volume. Since file systems are good at allocating contiguous logical blocks for a single file, or to files in the same directory, replicating logical segments is likely to identify and exploit most common access patterns.

[1]By being sensitive to data contents, semantically smart disks place a new requirement on file system traces to include user data blocks. However, the privacy concerns that such a campaign would encounter may be too difficult to overcome.

To implement access-driven diffusion, Alexander divides the logical address space into multiple segments, and, during normal operation, gathers information on the utilization of each segment. A background thread selects logical segments that remain "hot" for a certain number of consecutive *epochs* and diffuses a copy across the drives of the system. Subsequent reads and writes first go to these replicas, with background updates sent to the original blocks. The imap entry for the block indicates which copy is up to date. Clearly, the policy for deciding which segments to diffuse is quite simplistic in our prototype implementation. A more detailed analysis of the policy space for access-driven diffusion is left for future work.

The amount of disk space to allocate to performance-oriented replicas presents an important policy decision. The initial policy that Alexander implements is to reserve a certain minimum amount of space (specified by the system administrator) for these replicas, and then opportunistically use the free space available in the array for additional replication. This approach is similar to that used by AutoRAID for mirrored data [Wilkes et al. 1996], except that AutoRAID cannot identify data that is considered "dead" by the file system once written; in contrast, D-GRAID can use semantic knowledge to identify which blocks are free.

7.2 Live-Block Recovery

To implement live-block recovery, D-GRAID must understand which blocks are live. This knowledge must be correct in that no block that is live is considered dead, as that would lead to data loss. Alexander tracks this information by observing bitmap and data block traffic. Bitmap blocks tell us the liveness state of the file system that has been reflected to disk. However, due to reordering and delayed updates, it is not uncommon to observe a write to a data block whose corresponding bit has not yet been set in the data bitmap. To account for this, D-GRAID maintains a duplicate copy of all bitmap blocks, and whenever it sees a write to a block, sets the corresponding bit in the local copy of the bitmap. The duplicate copy is synchronized with the file system copy when the data bitmap block is written by the file system. This *conservative bitmap table* thus reflects a superset of all live blocks in the file system, and can be used to perform live-block recovery. Note that we assume the preallocation state of the bitmap will not be written to disk after a subsequent allocation; the locking in Linux and other modern systems already ensures this. Though this technique guarantees that a live block is never classified as dead, it is possible for the disk to consider a block live far longer than it actually is. This situation would arise, for example, if the file system writes deleted blocks to disk.

To implement live-block recovery, Alexander simply uses the conservative bitmap table to build a list of blocks which need to be restored. Alexander then proceeds through the list and copies all live data onto the hot spare.

7.3 Other Aspects of Alexander

There are a number of other aspects of the implementation that are required for a successful prototype. In this subsection, we briefly describe some of the key aspects.

7.3.1 *Physical Block Allocation.* The logical array of blocks exported by SCSI has the property that block numbers that are contiguous in the logical address space are mapped to contiguous physical locations on disk. This property empowers file systems to place data contiguously on disk simply by allocating contiguous logical blocks to the data. In traditional RAID, this property is straightforward to preserve. Because physical blocks are assigned in a round-robin fashion across disks, the contiguity guarantees continue to hold; the physical block to assign for a given logical block is a simple arithmetic calculation on the logical block number.

However in D-GRAID, deciding on the physical block to allocate for a newly written logical block is not straightforward; the decision depends on the file to which the logical block belongs, and its logical offset within the file. Because of fault-isolated placement, a set of contiguous logical blocks (e.g., those that belong to a single file) may all map to contiguous physical blocks on the same disk; thus, if a logical block L within that set is mapped to physical block P, block $L + k$ within the same set should be mapped to physical block $P + k$ in order to preserve contiguity expectations. However, at a larger granularity, since D-GRAID balances space utilization across files, the allocation policy should be different; for large values of k, block $L + k$ should map to physical block $P + (k/N)$ where N is the number of disks in the array. The choice of which of these policies to use requires estimates of file size which are quite dynamic.

Our prototype addresses this issue with a simple technique of space reservations. Alexander utilizes its knowledge of inodes and indirect blocks to get a priori estimates of the exact size of the entire file (or a large segment of the file, as in the case of indirect block). When it observes a new inode written that indicates a file of size b blocks, it reserves b contiguous blocks in the home site assigned for that file, so that when the actual logical blocks are written subsequently, the reserved space can be used. Note that since blocks are deferred until their inodes (or indirect blocks) are observed, a write to a new logical block will always have a prior reservation. Since inodes and indirect blocks are written only periodically (e.g., once every 5 s), the size information obtained from those writes is quite stable.

7.3.2 *Just-in-Time Commit.* Space reservations depend on the size information extracted from inode and indirect blocks. However, given that indirect block detection is fundamentally inaccurate, a misclassified indirect block could result in spurious reservations that hold up physical space. To prevent this, Alexander employs lazy allocation, where actual physical blocks are committed only when the corresponding logical block is written. The reservation still happens a priori, but these reservations are viewed as *soft* and the space is reclaimed if required.

7.3.3 *Interaction of Deferred Writes with* `sync`. Alexander defers disk writes of logical blocks for which it has not observed an owning inode. Such arbitrary deferral could potentially conflict with application-level expectations after a `sync` operation is issued; when a `sync` returns, the application expects

all data to be on disk. To preserve these semantics, D-GRAID handles inode and indirect block writes specially. D-GRAID does not return success on a write to an inode or indirect block until all deferred writes to blocks pointed to by that inode (or indirect) block have actually reached disk. Since the `sync` operation will not complete until the inode block write returns, all deferred writes are guaranteed to be complete before `sync` returns. The same argument extends for `fsync`, which will not return until all writes pertaining to the particular file complete. However, one weakness of this approach is that if the application performs an equivalent of `fdatasync` (i.e., which flushes only the data blocks to disk, and not metadata), the above technique would not preserve the expected semantics.

7.3.4 *Inconsistent Fault Behavior of Linux ext2.* One interesting issue that required a change from our design was the behavior of Linux ext2 under partial disk failure. When a process tries to read a data block that is unavailable, ext2 issues the read and returns an I/O failure to the process. When the block becomes available again (e.g., after recovery) and a process issues a read to it, ext2 will again issue the read, and everything works as expected. However, if a process tries to open a file whose inode is unavailable, ext2 marks the inode as "suspicious" and will never again issue an I/O request to the inode block, even if Alexander has recovered the block. To avoid a change to the file system and retain the ability to recover failed inodes, Alexander replicates inode blocks as it does namespace meta-data, instead of collocating them with the data blocks of a file.

7.3.5 *Persistence of Data Structures.* There are a number of structures that Alexander maintains, such as the imap, that must be reliably committed to disk and preferably, for good performance, buffered in a small amount of nonvolatile RAM. Note that since the NVRAM only needs to serve as a cache of actively accessed entries in these data structures, its space requirements can be kept at an acceptable level. Though our current prototype simply stores these data structures in memory, a complete implementation would require them to be backed persistently.

7.3.6 *Popular Directory Replication.* The most important component that is missing from the Alexander prototype is the decision on which "popular" (read-only) directories such as `/usr/bin` to replicate widely, and when to do so. Although Alexander contains the proper mechanisms to perform such replication, the policy space remains unexplored. However, our initial experience indicates that a simple approach based on monitoring frequency of inode access time updates may likely be effective. An alternative approach allows administrators to specify directories that should be treated in this manner.

7.4 Alexander the FAT

Overall, we were surprised by the many similarities we found in implementing D-GRAID underneath ext2 and VFAT. For example, VFAT also overloads data blocks, using them as either user data blocks or directories; hence Alexander

must defer classification of those blocks in a manner similar to the ext2 implementation. As can be expected, the implementation of most of the basic mechanisms in D-GRAID such as physical block allocation, allocation of home sites to files, and tracking replicas of critical blocks is shared across both versions of D-GRAID.

However, there were a few instances where the VFAT implementation of D-GRAID differed in interesting ways from the ext2 version. For example, the fact that all pointers of a file are located in the file allocation table made a number of aspects of D-GRAID much simpler to implement; in VFAT, there are no indirect pointers to worry about. When a new copy of a FAT block is written, the new version can be directly compared with the previous contents of the block to get accurate information on the blocks newly allocated or deleted. We also ran across the occasional odd behavior in the Linux implementation of VFAT. For example, Linux would write to disk certain blocks that were allocated but then freed, avoiding an obvious and common file system optimization. Because of this behavior of VFAT, our estimate of the set of live blocks will be a strict superset of the blocks that are actually live. Although this was more indicative of the untuned nature of the Linux implementation, it served as yet another indicator of how semantic disks must be wary of any assumptions they make about file system behavior.

8. EVALUATING ALEXANDER

We now present a performance evaluation of Alexander. We focus primarily on the Linux ext2 variant, but also include some baseline measurements of the VFAT system. We wish to answer the following questions:

—Does Alexander work correctly?
—What time overheads are introduced?
—How effective is access-driven diffusion?
—How fast is live-block recovery?
—What overall benefits can we expect from D-GRAID?
—How complex is the implementation?

8.1 Platform

The Alexander prototype is constructed as a software RAID driver in the Linux 2.2 kernel. File systems mount the pseudodevice and use it as if it were a normal disk. Our environment is excellent for understanding many of the issues that would be involved in the construction of a "real" hardware D-GRAID system; however, it is also limited in the following ways. First, and most importantly, Alexander runs on the same system as the host OS and applications, and thus there is interference due to competition for CPU and memory resources. Second, the performance characteristics of the microprocessor and memory system may be different than what is found within an actual RAID system. In the following experiments, we utilize a 550-MHz Pentium III and four 10K-rev/min IBM disks.

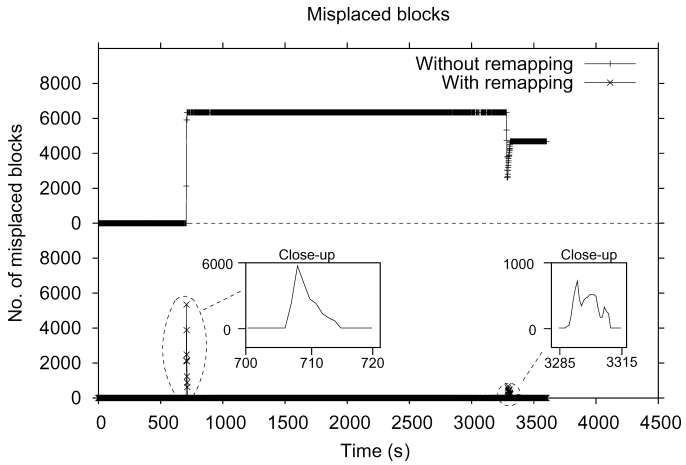

Fig. 5. Errors in placement. The figure plots the number of blocks wrongly laid out by Alexander over time, while running a busy hour of the HP Trace. The experiment was run over four disks, and the total number of blocks accessed in the trace was 418,000.

8.2 Does Alexander Work Correctly?

Alexander is more complex than simple RAID systems. To ensure that Alexander operates correctly, we have put the system through numerous stress tests, moving large amounts of data in and out of the system without problems. We have also extensively tested the corner cases of the system, pushing it into situations that are difficult to handle and making sure that the system degrades gracefully and recovers as expected. For example, we repeatedly crafted microbenchmarks to stress the mechanisms for detecting block reuse and for handling imperfect information about dynamically typed blocks. We have also constructed benchmarks that write user data blocks to disk that contain "worst-case" data, that is, data that appears to be valid directory entries or indirect pointers. In all cases, Alexander was able to detect which blocks were indirect blocks and move files and directories into their proper fault-isolated locations.

To verify that Alexander places blocks on the appropriate disk, we instrumented the file system to log block allocations. In addition, Alexander logs events of interest such as assignment of a home site for an inode, creation of a new mapping for a logical block, remapping of blocks to a different home site, and receipt of logical writes from the file system. To evaluate the behavior of Alexander on a certain workload, we ran the workload on Alexander, and obtained the time-ordered log of events that occurred at the file system and Alexander. We then processed this log off-line and looked for the number of blocks wrongly laid out at any given time.

We ran this test on a few hours of the HP Traces, and found that, in many of the hours we examined, the number of blocks that were misplaced even temporarily was quite low, often fewer than 10 blocks. We report detailed results for one such hour of the trace where we observed the greatest number of misplaced blocks, among the hours we examined. Figure 5 shows the results.

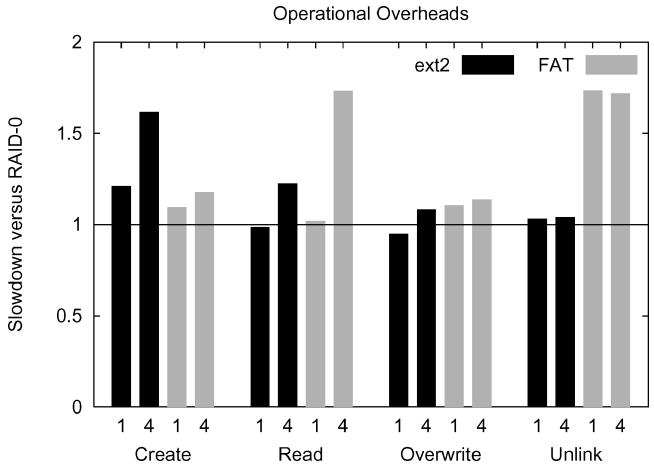

Fig. 6. Time overheads. The figure plots the time overheads observed on D-GRAID Level 0 versus RAID Level 0 across a series of microbenchmarks. The tests were run on one and four disk systems. In each experiment, 3000 operations were enacted (e.g., 3000 file creations), with each operation on a 64-kB file.

The figure has two parts. The bottom part shows the normal operation of Alexander, with the capability to react to block reuse by remapping (and copying over) blocks to the correct homesite. As the figure shows, Alexander is able to quickly detect wrongly placed blocks and remap them appropriately. Further, the number of such blocks misplaced temporarily is only about 1% of the total number of blocks accessed in the trace. The top part of the figure shows the number of misplaced blocks for the same experiment, but assuming that the remapping did not occur. As can be expected, those delinquent blocks remain misplaced. The dip toward the end of the trace occurs because some of the misplaced blocks were later assigned to a file in that home site itself (after a preceding delete), accidentally correcting the original misplacement.

8.3 What Time Overheads Are Introduced?

We now explore the time overheads that arise due to semantic inference. This primarily occurs when new blocks are written to the file system, such as during file creation. Figure 6 shows the performance of Alexander under a simple microbenchmark. As can be seen, allocating writes are slower due to the extra CPU cost involved in tracking fault-isolated placement. Reads and overwrites perform comparably with RAID-0. The high unlink times of D-GRAID on FAT is because FAT writes out data pertaining to deleted files, which have to be processed by D-GRAID as if it were newly allocated data. Given that the implementation is untuned and the infrastructure suffers from CPU and memory contention with the host, we believe that these are worst-case estimates of the overheads.

Another cost of D-GRAID that we explored was the overhead of metadata replication. For this purpose, we chose Postmark [Katcher 1997], a

Table II. Performance on Postmark. (The table compares the performance of D-GRAID Level 0 with RAID-0 on the Postmark benchmark. Each row marked D-GRAID indicates a specific level of metadata replication. The first column reports the benchmark run-time and the second column shows the number of disk writes incurred. The third column shows the number of disk writes that were to metadata blocks, and the fourth column indicates the number of unique metadata blocks that are written. The experiment was run over four disks)

		Blocks written		
	Run-time(s)	Total	Meta-data	Unique
RAID-0	69.25	101,297	—	—
D-GRAID$_1$	61.57	93,981	5962	1599
D-GRAID$_2$	66.50	99,416	9954	3198
D-GRAID$_3$	73.50	101,559	16,976	4797
D-GRAID$_4$	78.79	113,222	23,646	6396

meta-data-intensive file system benchmark. We slightly modified Postmark to perform a `sync` before the deletion phase, so that all meta-data writes were accounted for, making it a pessimistic evaluation of the costs. Table II shows the performance of Alexander under various degrees of meta-data replication. As can be seen from the table, synchronous replication of meta-data blocks had a significant effect on performance for meta-data-intensive workloads (the file sizes in Postmark range from 512 bytes to 10 kB). Note that Alexander performed better than default RAID-0 for lower degrees of replication because of better physical block allocation; since ext2 looks for a contiguous free chunk of eight blocks to allocate a new file, its layout is suboptimal for small files, since it does not pack them together.

The table also shows the number of disk writes incurred during the course of the benchmark. The percentage of extra disk writes roughly accounts for the difference in performance between different replication levels, and these extra writes were mostly to meta-data blocks. However, when we counted the number of unique physical writes to meta-data blocks, the absolute difference between different replication levels was small. This suggests that lazy propagation of updates to meta-data block replicas, perhaps during idle time or using free-block scheduling, can greatly reduce the performance difference, at the cost of added complexity. For example, with lazy update propagation (i.e., if the replicas were updated only once), D-GRAID$_4$ would incur only about 4% extra disk writes.

We also played back a portion of the HP traces for 20 min against a standard RAID-0 system and D-GRAID over four disks. The playback engine issued requests at the times specified in the trace, with an optional speedup factor; a speedup of 2× implies the idle time between requests was reduced by a factor of 2. With speedup factors of 1× and 2×, D-GRAID delivered the same per-second operation throughput as RAID-0, utilizing idle time in the trace to hide its extra CPU overhead. However, with a scaling factor of 3×, the operation throughput lagged slightly behind, with D-GRAID showing a slowdown of up to 19.2% during the first one-third of the trace execution, after which it caught up due to idle time.

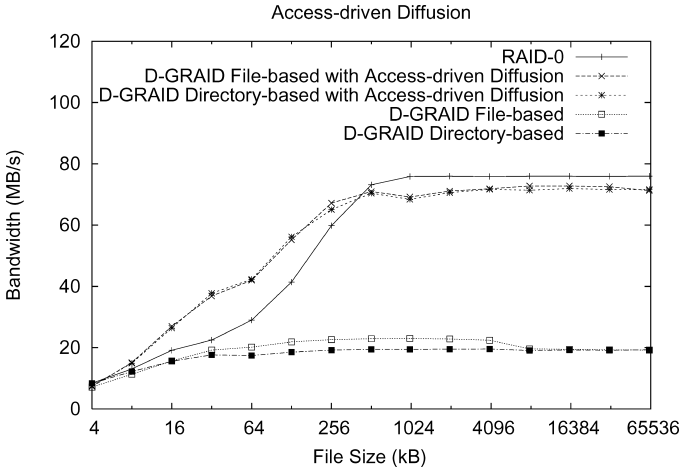

Fig. 7. Access-driven diffusion. The figure presents the performance of D-GRAID Level 0 and standard RAID-0 under a sequential workload. In each experiment, a number of files of size x are read sequentially, with the total volume of data fixed at 64 MB. D-GRAID performs better for smaller files due to better physical block layout.

8.4 How Effective Is Access-Driven Diffusion?

We now show the benefits of access-driven diffusion. In each trial of this experiment, we performed a set of sequential file reads, over files of increasing size. We compared standard RAID-0 striping to D-GRAID with and without access-driven diffusion. Figure 7 shows the results of the experiment.

As we can see from the figure, without access-driven diffusion, sequential access to larger files ran at the rate of a single disk in the system, and thus did not benefit from the potential parallelism. With access-driven diffusion, performance was much improved, as reads were directed to the diffused copies across all of the disks in the system. Note that, in the latter case, we arranged for the files to be already diffused before the start of the experiment, by reading them a certain threshold number of times. Investigating more sophisticated policies for when to initiate access-driven diffusion is left for future work.

8.5 How Fast Is Live-Block Recovery?

We now explore the potential improvement seen with live-block recovery. Figure 8 presents the recovery time of D-GRAID while varying the amount of live file system data.

The figure plots two lines: worst-case and best-case live-block recovery. In the worst case, live data is spread throughout the disk, whereas in the best case it is compacted into a single portion of the volume. From the graph, we can see that live-block recovery was successful in reducing recovery time, particularly when a disk was less than half full. Note also the difference between worst-case and best-case times; the difference suggests that periodic disk reorganization [Ruemmler and Wilkes 1991] could be used to speed recovery, by moving all live data to a localized portion.

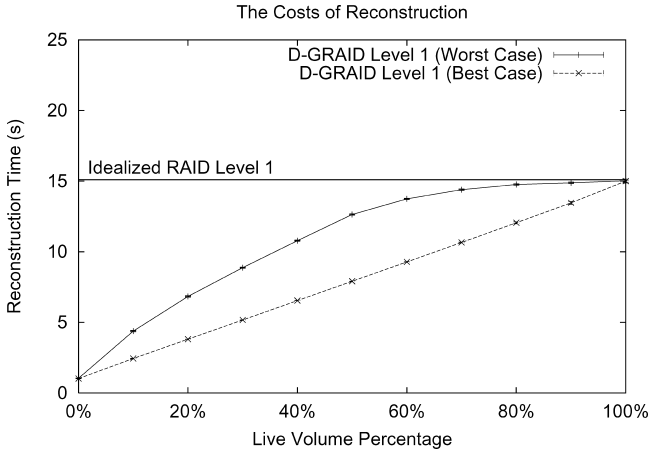

Fig. 8. Live-block recovery. The figure shows the time to recover a failed disk onto a hot spare in a D-GRAID Level 1 (mirrored) system using live-block recovery. Two lines for D-GRAID are plotted: in the worst case, live data is spread across the entire 300-MB volume, whereas in the best case it is compacted into the smallest contiguous space possible. Also plotted is the recovery time of an idealized RAID Level 1.

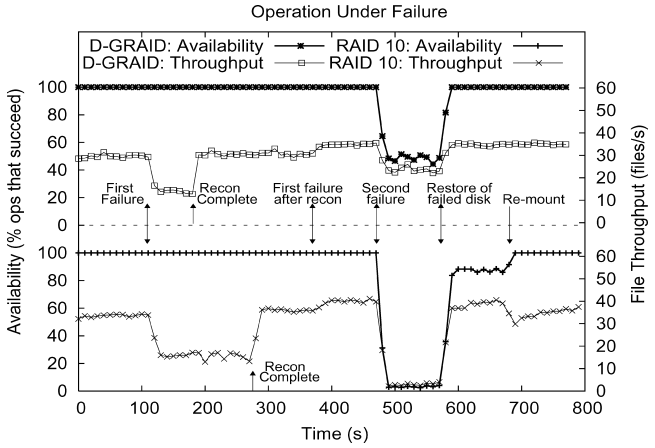

Fig. 9. Availability profile. The figure shows the operation of D-GRAID Level 1 and RAID 10 under failures. The 3-GB array consisted of four data disks and one hot spare. After the first failure, data was reconstructed onto the hot spare, D-GRAID recovering much faster than RAID 10. When two more failures occur, RAID 10 loses almost all files, while D-GRAID continued to serve 50% of its files. The workload consisted of read-modify-writes of 32-kB files randomly picked froma 1.3-GB working set.

8.6 What Overall Benefits Can We Expect from D-GRAID?

We next demonstrate the improved availability of Alexander under failures. Figure 9 shows the availability and performance observed by a process randomly accessing whole 32-kB files, running above D-GRAID and RAID-10. To ensure a fair comparison, both D-GRAID and RAID-10 limited their reconstruction rate to 10 MB/s.

Table III. Code Size for Alexander Implementation (The number of lines of code needed to implement Alexander is shown. The first column shows the number of semicolons and the second column shows the total number of lines, including white spaces and comments.)

	Semicolons	Total
D-GRAID generic		
Setup + fault-isolated placement	1726	3557
Physical block allocation	322	678
Access driven diffusion	108	238
Mirroring + live block recovery	248	511
Internal memory management	182	406
Hashtable/Avl tree	724	1706
File system specific		
SDS Inferencing: ext2	1252	2836
SDS Inferencing: VFAT	630	1132
Total	5192	11604

As the figure shows, reconstruction of the 3-GB volume with 1.3-GB live data completed much faster (68 s) in D-GRAID compared with RAID-10 (160 s). Also, when the extra second failure occured, the availability of RAID-10 dropped to near zero, while D-GRAID continued with about 50 % availability. Surprisingly, after restore, RAID-10 still failed on certain files; this is because Linux does not retry inode blocks once they fail. A remount is required before RAID-10 returns to full availability.

8.7 How Complex Is the Implementation?

We briefly quantify the implementation complexity of Alexander. Table III shows the number of C statements required to implement the different components of Alexander. From the table, we can see that the core file system inferencing module for ext2 requires only about 1200 lines of code (counted with number of semicolons), and the core mechanisms of D-GRAID contribute to about 2000 lines of code. The rest is spent on a hash table, AVL tree, and wrappers for memory management. Compared to the tens of thousands of lines of code already comprising modern array firmware, we believe that the added complexity of D-GRAID is not that significant. Being an academic prototype, these complexity numbers could be a slight underestimate of what would be required for a production quality implementation; thus, this analysis is only intended to be an approximate estimate.

9. D-GRAID LEVELS

Much of the discussion so far has focused on implementing D-GRAID over a storage system with no redundancy for user data (i.e., RAID-0), or over a mirrored storage system (i.e., RAID-10). However, as mentioned before, the layout *mechanisms* in D-GRAID are orthogonal to the underlying redundancy scheme. In this section, we formalize the different *levels* of D-GRAID, corresponding to the popular traditional RAID levels. We also present certain custom *policies* for each D-GRAID level that are tailored to the underlying redundancy mechanism.

Note that, in contrast to traditional RAID levels, the levels of D-GRAID differ only in the type of redundancy for normal user data; system meta-data is always maintained in RAID-1 with a certain configured replication degree.

9.1 D-GRAID-0: No Redundancy

This is the simplest D-GRAID level where no redundancy mechanism is employed for normal user data. Thus, even a single disk failure results in data loss. In contrast to traditional RAID-0 where a single disk failure results in complete data loss, D-GRAID-0 ensures proportional data availability under failure. Figure 10(a) shows the D-GRAID-0 configuration.

Because of the absence of redundancy for normal data, the additional storage required for access-driven diffusion in D-GRAID-0 needs to come from a separate *performance reserve*, as described in Section 7. This reserve can be fixed to be a certain percentage (e.g., 10% of the storage volume size) or can be tunable by the administrator. Tuning this parameter provides the administrator control over the tradeoff between performance and storage efficiency. One issue with changing the size of the performance reserve dynamically is that file systems may not be equipped to deal with a variable volume size. This limitation can be addressed by a simple technique: the administrator creates a file in the file system with a certain reserved name (e.g., */.diffuse*). The size of this file implicitly conveys to D-GRAID the size of its performance reserve. Since the file system will not use the blocks assigned to this reserved file to any other file, D-GRAID is free to use this storage space. When the file system runs short of storage, the administrator can prune the size of this special file, thus dynamically reducing the size of the performance reserve.

9.2 D-GRAID-10: Mirroring

A mirrored D-GRAID system stripes data across multiple mirrored pairs, similar to RAID-10. Note that D-GRAID is not meaningful in a storage system comprised of a single mirrored pair (i.e., RAID-1) because such a system fundamentally has no partial failure mode. The access-driven diffusion policy in D-GRAID-10 is quite similar to D-GRAID-0 where a dynamic performance reserve is used to hold diffused copies; Figure 10(b) depicts this configuration. Note that the diffused copies are not mirrored; thus D-GRAID-10 requires only half the percentage of space that D-GRAID-0 requires, in order to achieve the same level of diffusion.

A slight variant of D-GRAID-10 can make access-driven diffusion much more effective, though at the cost of a slight degradation in reliability. Instead of the disks in a mirrored pair being physical mirrors as discussed above, we could employ *logical* mirroring, where we just impose that each logical disk block has two copies in two different disks. With such a relaxed definition, D-GRAID could store one copy of a file in the traditional striped fashion, while the other copy of the file is stored in fault-isolated fashion. Figure 10(c) depicts this configuration. Each file has a fault-isolated copy laid out in a single disk, and another copy striped across all the other disks, so that a single disk failure will not result in any data loss. Such logical mirroring of data achieves the benefits of

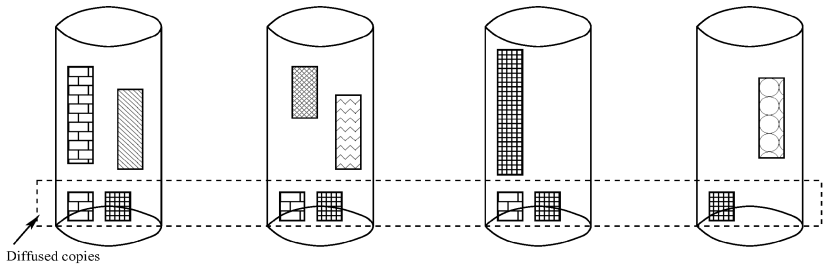

(a) DGRAID − 0

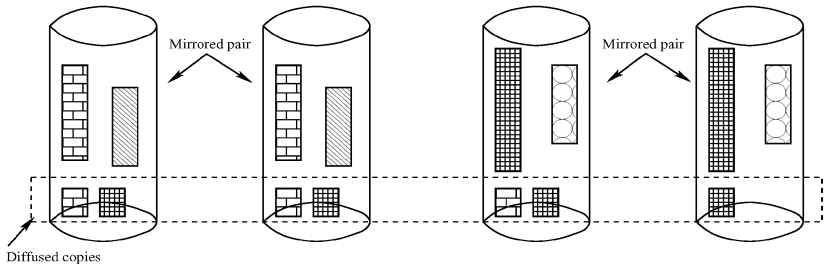

(b) DGRAID − 1 (Physical)

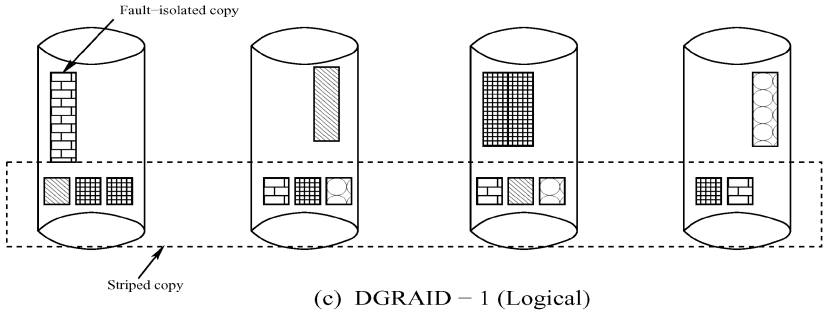

(c) DGRAID − 1 (Logical)

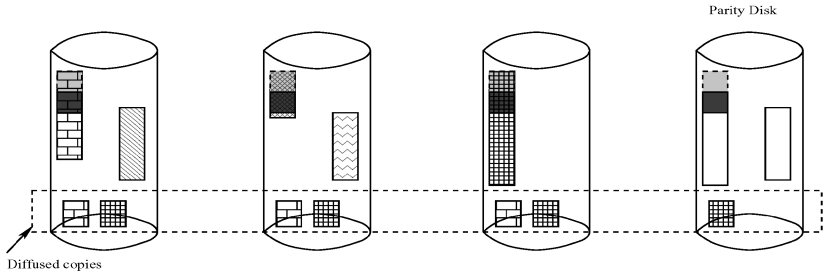

(d) DGRAID − 5

Fig. 10. D-GRAID levels. The figures depict the data layout of D-GRAID under various redundancy schemes. Each style of shading represents a different file. In the D-GRAID-5 figure, the color of shading indicates a physical RAID-5 stripe. The diffusion segments and the striped region in D-GRAID-1(Logical) are indicated as separate regions of the disk for simplicity; in practice, they will be interleaved with the fault-isolated copies.

fault-isolated placement with almost no impact on performance, because parallelism is still available out of the striped copies. Note that in such a scenario no extra space is required for access-driven diffusion.

Although the above variant of D-GRAID-10 improves performance by more efficient access-driven diffusion, it reduces reliability compared to a traditional D-GRAID-10. In a traditional D-GRAID-10 (i.e., physical mirroring), after a single disk failure, only the failure of its mirror disk will lead to loss of data. However, in logical mirroring, the second failure always results in loss of data, though proportionally, irrespective of which disk incurred the failure.

9.3 D-GRAID-5: Parity

D-GRAID-5 is the counterpart of traditional RAID-5; redundancy for user data is maintained in the form of parity encoding on a small number of disks (usually 1), resulting in better space efficiency. While it may appear that the fine grained block-level striping that is fundamental to RAID-5 would be in conflict with the fault isolated placement in D-GRAID, these techniques are quite orthogonal. The fine-grained striping required for RAID-5 occurs at the *physical* level, across actual physical disk blocks, while fault-isolated placement is just a logical assignment of files onto those physical blocks. Thus, D-GRAID-5 would still maintain the invariant that the kth parity block is the XOR of the kth block in every disk; the only difference is that the kth block in each disk would contain data pertaining to a different file in D-GRAID, while in RAID, they would usually be part of the same file. This configuration is shown in Figure 10(d), where blocks belonging to the same physical RAID-5 stripe are shaded with the same color.

However, fault-isolated placement with RAID-5 like redundancy leads to a performance issue. Since blocks within a RAID-5 stripe are no longer part of a single file (and thus not logically related), full stripe writes become uncommon. Thus with the block allocation policies described so far, most writes will be to partial stripes; such *small writes* have the well known performance problem of requiring four disk operations for every block written [Patterson et al. 1988].

To address the small write problem in D-GRAID-5, we need a customized block allocation policy. While the allocation policies described in Section 7 are targeted at preserving the logical contiguity perceived by the file system, D-GRAID-5 requires a policy that minimizes the impact of small writes. One example of such a policy is log-structured allocation [Rosenblum and Ousterhout 1992; Wilkes et al. 1996], where blocks are not written in place, but allocated from empty *segments*, invalidating the old locations.

With such log structured allocation, D-GRAID-5 would simply divide each disk into multiple segments; at any given time, D-GRAID-5 would operate on a *segment stripe*, which comprises of the kth segment in each disk. When a write arrives, the fault isolation module of D-GRAID-5 would decide which disk the block needs to be laid out in, and then would allocate the tail physical block of the corresponding segment to that logical block. Considering that in a typical workload, writes are spread across multiple files, and given that D-GRAID balances space utilization across disks, it is most likely that writes

to such multiple files are spread across different segments within the current segment stripe, thus resulting in full stripe writes. Note however that for this technique to be effective, the *log cleaner* should coordinate cleaning across the entire set of disks, so that the set of freed segments comprise full segment stripes.

9.4 Summary

In summary, we find that the basic layout techniques in D-GRAID are orthogonal to the underlying redundancy mechanism. By building on top of any physical redundancy scheme, D-GRAID strictly improves the availability of the storage array. However, custom policies (e.g., for access-driven diffusion, physical block allocation, etc.) often make D-GRAID more effective for a given redundancy mechanism.

10. DISCUSSION: THE IMPACT OF BEING WRONG

As described in Section 7, there is a fair amount of complexity in identifying the logical file to which a block belongs, in order to place it in the correct home site for graceful degradation. An interesting question that arises in the light of such complexity is: what happens if D-GRAID makes a wrong inference? For example, what happens if D-GRAID permanently associates a block with the wrong file, and thus places it in the wrong home site? Such incorrect inferences affect different parts of the D-GRAID design differently.

The graceful degradation component of D-GRAID is quite robust to incorrect inferences; an incorrect association of a block to the wrong file would only affect fault isolation, and not impact correctness. Even if D-GRAID miscalculates a large fraction of its associations, the reliability of the resulting storage layout will still be strictly better than the corresponding traditional RAID level. This is because D-GRAID builds on top of existing RAID redundancy. An incorrect association may lead to a layout that is not completely fault isolated, but such a layout will still exhibit better fault isolation compared to traditional RAID. Thus even in the face of incorrect inference, the storage system correctness is not affected, thus making D-GRAID an ideal candidate to make aggressive use of such semantic information.

In contrast, the live block recovery component of D-GRAID does depend on semantic information for correctness. Although it requires only a conservative estimate of the set of live blocks in the volume, D-GRAID requires this estimate to be *strictly* conservative; a live block should never be inferred to be dead, since that could lead to loss of data. However, as described in Section 7, tracking such block liveness information conservatively is quite simple, and thus is straightforward to realize.

Thus, D-GRAID requires accuracy only for a very simple piece of semantic information for implementing fast recovery. Much of the design and complexity of D-GRAID is related to fault isolation for graceful degradation; this component is much more robust to incorrect inference, and cannot be "wrong" in any bad way.

11. CONCLUSIONS

> "A robust system is one that continues to operate (nearly) correctly in the presence of some class of errors." *Robert Hagmann [Hagmann 1987]*

D-GRAID turns the simple binary failure model found in most storage systems into a continuum, increasing the availability of storage by continuing operation under partial failure and quickly restoring live data after a failure does occur. In this article, we have shown the potential benefits of D-GRAID, established the limits of semantic knowledge, and shown how a successful D-GRAID implementation can be achieved despite these limits. Through simulation and the evaluation of a prototype implementation, we have found that D-GRAID can be built underneath a standard block-based interface, without any file system modification, and that it delivers graceful degradation and live-block recovery, and, through access-driven diffusion, good performance.

We conclude with a discussions of the lessons we have learned in the process of implementing D-GRAID:

—*Limited knowledge within the disk does not imply limited functionality.* One of the main contributions of this article is a demonstration of both the limits of semantic knowledge, as well as the "proof" via implementation that despite such limitations, interesting functionality can be built inside of a semantically smart disk system. We believe any semantic disk system must be careful in its assumptions about file system behavior, and hope that our work can guide others who pursue a similar course.

—*Semantically smart disks would be easier to build with some help from above.* Because of the way file systems reorder, delay, and hide operations from disks, reverse engineering exactly what they are doing at the SCSI level is difficult. We believe that small modifications to file systems could substantially lessen this difficulty. For example, if the file system could inform the disk whenever it believes the file system structures are in a consistent on-disk state, many of the challenges in the disk would be lessened. This is one example of many small alterations that could ease the burden of semantic disk development.

—*Semantically smart disks stress file systems in unexpected ways.* File systems were not built to operate on top of disks that behave as D-GRAID does; specifically, they may not behave particularly well when part of a volume address space becomes unavailable. Perhaps because of its heritage as an OS for inexpensive hardware, Linux file systems handle unexpected conditions fairly well. However, the exact model for dealing with failure is inconsistent: data blocks could be missing and then reappear, but the same is not true for inodes. As semantically smart disks push new functionality into storage, file systems may potentially need to evolve to accommodate them.

ACKNOWLEDGMENTS

We would like to thank Anurag Acharya, Erik Riedel, Yasushi Saito, John Bent, Nathan Burnett, Timothy Denehy, Brian Forney, Florentina Popovici, and Lakshmi Bairavasundaram for their insightful comments on earlier drafts of the article, and Jack Harwood for helpful discussions. We also would like to

thank Richard Golding for his excellent shepherding of an earlier version of the article, and the anonymous reviewers for their thoughtful suggestions, many of which have greatly improved the content of this article. Finally, we thank the Computer Systems Lab for providing a terrific environment for computer science research.

REFERENCES

ACHARYA, A., UYSAL, M., AND SALTZ, J. 1998. Active disks: Programming model, algorithms and evaluation. In *Proceedings of the 8th International Conference on Architectural Support for Programming Languages and Operating Systems* (ASPLOS VIII, San Jose, CA).

ALVAREZ, G. A., BURKHARD, W. A., AND CRISTIAN, F. 1997. Tolerating multiple failures in RAID architectures with optimal storage and uniform declustering. In *Proceedings of the 24th Annual International Symposium on Computer Architecture* (ISCA '97, Denver, CO).

ANDERSON, D., CHASE, J., AND VAHDAT, A. 2002. Interposed request routing for scalable network storage. *ACM Trans. Comput. Syst. 20*, 1 (Feb.), 25–48.

BITTON, D. AND GRAY, J. 1988. Disk shadowing. In *Proceedings of the 14th International Conference on Very Large Data Bases* (VLDB 14, Los Angeles, CA). 331–338.

BOEHM, H. AND WEISER, M. 1988. Garbage collection in an uncooperative environment. *Softw.—Pract. Exper. 18*, 9 (Sep.), 807–820.

BURKHARD, W. AND MENON, J. 1993. Disk array storage system reliability. In *Proceedings of the 23rd International Symposium on Fault-Tolerant Computing* (FTCS-23, Toulouse, France). 432–441.

CHAPIN, J., ROSENBLUM, M., DEVINE, S., LAHIRI, T., TEODOSIU, D., AND GUPTA, A. 1995. Hive: Fault containment for shared-memory multiprocessors. In *Proceedings of the 15th ACM Symposium on Operating Systems Principles* (SOSP '95, Copper Mountain Resort, CO).

CHEN, P. M., LEE, E. K., GIBSON, G. A., KATZ, R. H., AND PATTERSON, D. A. 1994. RAID: High-performance, reliable secondary storage. *ACM Comput. Surv. 26*, 2 (June), 145–185.

DENEHY, T. E., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2002. Bridging the information gap in storage protocol stacks. In *Proceedings of the USENIX Annual Technical Conference* (USENIX '02, Monterey, CA).

DOWSE, I. AND MALONE, D. 2002. Recent filesystem optimisations on FreeBSD. In *Proceedings of the USENIX Annual Technical Conference* (FREENIX Track, Monterey, CA).

EMC CORPORATION. 2002. Symmetrix Enterprise Information Storage Systems. EMC Corporation, Hopkinton, MA. Web site: `http://www.emc.com`.

ENGLISH, R. M. AND STEPANOV, A. A. 1992. Loge: A self-organizing disk controller. In *Proceedings of the USENIX Winter Technical Conference* (USENIX Winter '92, San Francisco, CA).

GANGER, G. R. 2001. Blurring the line between oses and storage devices. Tech. rep. CMU-CS-01-166. Carnegie Mellon University, Pittsburgh, PA.

GANGER, G. R., MCKUSICK, M. K., SOULES, C. A., AND PATT, Y. N. 2000. Soft updates: A solution to the metadata update problem in file systems. *ACM Trans. Comput. Syst. 18*, 2 (May), 127–153.

GANGER, G. R., WORTHINGTON, B. L., HOU, R. Y., AND PATT, Y. N. 1993. Disk subsystem load balancing: Disk striping vs. conventional data placement. In *HICSS '93*.

GIBSON, G. A., NAGLE, D. F., AMIRI, K., BUTLER, J., CHANG, F. W., GOBIOFF, H., HARDIN, C., RIEDEL, E., ROCHBERG, D., AND ZELENKA, J. 1998. A cost-effective, high-bandwidth storage architecture. In *Proceedings of the 8th International Conference on Architectural Support for Programming Languages and Operating Systems* (ASPLOS VIII, San Jose, CA).

GRAY, J. 1987. Why do computers stop and what can we do about it? In *Proceedings of the 6th International Conference on Reliability and Distributed Databases*.

GRAY, J., HORST, B., AND WALKER, M. 1990. Parity striping of disc arrays: Low-cost reliable storage with acceptable throughput. In *Proceedings of the 16th International Conference on Very Large Data Bases* (VLDB 16, Brisbane, Australia). 148–159.

GRIBBLE, S. D. 2001. Robustness in complex systems. In *Proceedings of the Eighth Workshop on Hot Topics in Operating Systems* (HotOS VIII, Schloss Elmau, Germany).

HAGMANN, R. 1987. Reimplementing the Cedar file system using logging and group commit. In *Proceedings of the 11th ACM Symposium on Operating Systems Principles* (SOSP '87, Austin, Texas).

HOLLAND, M., GIBSON, G., AND SIEWIOREK, D. 1993. Fast, on-line failure recovery in redundant disk arrays. In *Proceedings of the 23rd International Symposium on Fault-Tolerant Computing* (FTCS-23, Toulouse, France).

HSIAO, H.-I. AND DEWITT, D. 1990. Chained declustering: A new availability strategy for multiprocessor database machines. In *Proceedings of the 6th International Data Engineering Conference*.

IBM. 2001. ServeRAID—recovering from multiple disk failures. Web site: `http://www.pc.ibm.com/qtechinfo/MIGR-39144.html`.

JI, M., FELTEN, E., WANG, R., AND SINGH, J. P. 2000. Archipelago: An island-based file system for highly available and scalable Internet services. In *Proceedings of the 4th USENIX Windows Symposium*.

KATCHER, J. 1997. PostMark: A new file system benchmark. Tech. rep. TR-3022, Network Appliance Inc., Sunnyvale, CA. Web site: `http://www.netapp.com`.

KEETON, K. AND WILKES, J. 2002. Automating data dependability. In *Proceedings of the 10th ACM-SIGOPS European Workshop*. (Saint-Emilion, France). 93–100.

KISTLER, J. AND SATYANARAYANAN, M. 1992. Disconnected operation in the Coda file system. *ACM Trans. Comput. Syst. 10*, 1 (Feb.), 3–25.

MCKUSICK, M. K., JOY, W. N., LEFFLER, S. J., AND FABRY, R. S. 1984. A fast file system for UNIX. *ACM Trans. Comput. Syst. 2*, 3 (Aug.), 181–197.

MENON, J. AND MATTSON, D. 1992. Comparison of sparing alternatives for disk arrays. In *ISCA '92*. (Gold Coast, Australia).

MICROSOFT CORPORATION. 2000. Web site: `http://www.microsoft.com/hwdev/`.

ORJI, C. U. AND SOLWORTH, J. A. 1993. Doubly distorted mirrors. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data* (SIGMOD '93, Washington, DC).

PARK, A. AND BALASUBRAMANIAN, K. 1986. Providing fault tolerance in parallel secondary storage systems. Tech. rep. CS-TR-057-86. Princeton, University, Princeton, NJ.

PATTERSON, D., GIBSON, G., AND KATZ, R. 1988. A case for redundant arrays of inexpensive disks (RAID). In *Proceedings of the 1988 ACM SIGMOD Conference on the Management of Data* (SIGMOD '88, Chicago, IL).

PATTERSON, D. A. 2002. Availability and maintainability $\gg$ performance: New focus for a new century. Key note speech at FAST '02.

POPEK, G., WALKER, B., CHOW, J., EDWARDS, D., KLINE, C., RUDISIN, G., AND THIEL, G. 1981. LOCUS: A network transparent, high reliability distributed system. In *Proceedings of the 8th ACM Symposium on Operating Systems Principles* (SOSP '81, Pacific Grove, CA).

REDDY, A. L. N. AND BANERJEE, P. 1991. Gracefully degradable disk arrays. In *Proceedings of the 21st International Symposium on Fault-Tolerant Computing* (FTCS-21, Montreal, P.Q. Canada). 401–408.

RIEDEL, E., GIBSON, G., AND FALOUTSOS, C. 1998. Active storage for large-scale data mining and multimedia. In *Proceedings of the 24th International Conference on Very Large Databases* (VLDB 24, New York, NY).

RIEDEL, E., KALLAHALLA, M., AND SWAMINATHAN, R. 2002. A framework for evaluating storage system security. In *Proceedings of the 1st USENIX Symposium on File and Storage Technologies* (FAST '02, Monterey, CA). 14–29.

ROSENBLUM, M. AND OUSTERHOUT, J. 1992. The design and implementation of a log-structured file system. *ACM Trans. Comput. Syst. 10*, 1 (Feb.), 26–52.

ROWSTRON, A. AND DRUSCHEL, P. 2001. 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* (SOSP '01, Banff, Alto., Canada).

RUEMMLER, C. AND WILKES, J. 1991. Disk shuffling. Tech. rep. HPL-91-156. Hewlett Packard Laboratories, Palo Alto, CA.

SAITO, Y., KARAMANOLIS, C., KARLSSON, M., AND MAHALINGAM, M. 2002. Taming aggressive replication in the Pangaea wide-area file system. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation* (OSDI '02, Boston, MA).

SAVAGE, S. AND WILKES, J. 1996. AFRAID—a frequently redundant array of independent disks. In *Proceedings of the USENIX Annual Technical Conference* (USENIX '96, San Diego, CA). 27–39.

SIVATHANU, M., PRABHAKARAN, V., POPOVICI, F. I., DENEHY, T. E., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. 2003. Semantically-smart disk systems. In *FAST '03* (San Francisco, CA). 73–88.

TS'O, T. AND TWEEDIE, S. 2002. Future directions for the Ext2/3 filesystem. In *Proceedings of the USENIX Annual Technical Conference* (FREENIX Track, Monterey, CA).

WANG, R., ANDERSON, T. E., AND PATTERSON, D. A. 1999. Virtual log-based file systems for a programmable disk. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation* (OSDI '99, New Orleans, LA).

WILKES, J., GOLDING, R., STAELIN, C., AND SULLIVAN, T. 1996. The HP AutoRAID hierarchical storage system. *ACM Trans. Comput. Syst. 14*, 1 (Feb.), 108–136.

WOLF, J. L. 1989. The placement optimization problem: A practical solution to the disk file assignment problem. In *Proceedings of the 1989 ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems* (SIGMETRICS '89, Berkeley, CA). 1–10.

Received August 2004; revised August 2004; accepted September 2004