

A Study of Linux File System Evolution

Lanyue Lu, Andrea C. Arpaci-Dusseau*, Remzi H. Arpaci-Dusseau*, Shan Lu*
Computer Sciences Department, University of Wisconsin, Madison

* Joint senior authors, listed alphabetically

Abstract

We conduct a comprehensive study of file-system code evolution. By analyzing eight years of Linux file-system changes across 5079 patches, we derive numerous new (and sometimes surprising) insights into the file-system development process; our results should be useful for both the development of file systems themselves as well as the improvement of bug-finding tools.

1 Introduction

Open-source local file systems, such as Linux Ext4 [31], XFS [46], and Btrfs [30], remain a critical component in the world of modern storage. For example, many recent distributed file systems, such as Google GFS [17] and Hadoop DFS [43], all replicate data objects (and associated metadata) across local file systems. On smart phones, most user data is managed by a local file system; for example, Google Android phones use Ext4 [2, 23] and Apple’s iOS devices use HFSX [34]. Finally, many desktop users still do not backup their data regularly [21, 29]; in this case, the local file system clearly plays a critical role as sole manager of user data.

Open-source local file systems remain a moving target. Developed by different teams with different goals, these file systems evolve rapidly to add new features, fix bugs, and improve performance and reliability, as one might expect in the open-source community [38]. Major new file systems are introduced every few years [12, 30, 32, 39, 46]; with recent technology changes (e.g., Flash [11, 18]), we can expect even more flux in this domain.

However, despite all the activity in local file system development, there is little *quantitative* understanding of their code bases. For example, where does the complexity of such systems lie? What types of bugs are common? Which performance features exist? Which reliability features are utilized? These questions are important to answer for different communities: for developers, so that they can improve current designs and implementations and create better systems; for tool builders, so that they can improve their tools to match reality (e.g., by finding the types of bugs that plague existing systems).

One way to garner insight into these questions is to study the artifacts themselves. Compared with proprietary software, open source projects provide a rich resource for source code and patch analysis. The fact that every version of Linux is available online, including a detailed set

of patches which describe how one version transforms to the next, enables us to carefully analyze how file systems have changed over time. A new type of “systems software archeology” is now possible.

In this paper, we perform the first comprehensive study of the evolution of Linux file systems, focusing on six major and important ones: Ext3 [47], Ext4 [31], XFS [46], Btrfs [30], ReiserFS [13], and JFS [10]. These file systems represent diverse features, designs, implementations and even groups of developers. We examine every file-system patch in the Linux 2.6 series over a period of eight years including 5079 patches. By carefully studying each patch to understand its intention, and then labeling the patch accordingly along numerous important axes, we can gain deep quantitative insight into the file-system development process. We can then answer questions such as “what are most patches for?”, “what types of bugs are common?”, and in general gain a new level of insight into the common approaches and issues that underlie current file-system development and maintenance.

We make the following high-level observations (§3). A large number of patches (nearly 50%) are maintenance patches, reflecting the constant refactoring work needed to keep code simple and maintainable. The remaining dominant category is bugs (just under 40%, about 1800 bugs), showing how much effort is required to slowly inch towards a “correct” implementation; perhaps this hard labor explains why some have found that the quality of open source projects is better than the proprietary software average [1]. Interestingly, the number of bugs does not die down over time (even for stable file systems), rather ebbing and flowing over time.

Breaking down the bug category further (§4), we find that semantic bugs, which require an understanding of file-system semantics to find or fix, are the dominant bug category (over 50% of all bugs). These types of bugs are vexing, as most of them are hard to detect via generic bug detection tools [9, 35]; more complex model checking [52] or formal specification [24] may be needed. Concurrency bugs are the next most common (about 20% of bugs), more prevalent than in user-level software [26, 42, 45]. Within this group, atomicity violations and deadlocks dominate. Kernel deadlocks are common (many caused by incorrectly using blocking kernel functions), hinting that recent research [22, 49] might be needed in-kernel. The remaining bugs are split relatively

evenly across memory bugs and improper error-code handling. In the memory bug category, memory leaks and null-pointer dereferences are common; in the error-code category, most bugs simply drop errors completely [19].

We also categorize bugs along other axes to gain further insight. For example, when broken down by consequence, we find that most of the bugs we studied lead to crashes or corruption, and hence are quite serious; this result holds across semantic, concurrency, memory, and error code bugs. When categorized by data structure, we find that B-trees, present in many file systems for scalability, have relatively few bugs per line of code. When classified by whether bugs occur on normal or failure-handling paths, we make the following important discovery: nearly 40% of all bugs occur on failure-handling paths. File systems, when trying to react to a failed memory allocation, I/O error, or some other unexpected condition, are highly likely to make further mistakes, such as incorrect state updates and missing resource releases. These mistakes can lead to corruption, crashes, deadlocks and leaks. Future system designs need better tool or language support to make these rarely-executed failure paths correct.

Finally, while bug patches comprise most of our study, performance and reliability patches are also prevalent, accounting for 8% and 7% of patches respectively (§5). The performance techniques used are relatively common and widespread (e.g., removing an unnecessary I/O, or downgrading a write lock to a read lock). About a quarter of performance patches reduce synchronization overheads; thus, while correctness is important, performance likely justifies the use of more complicated and time saving synchronization schemes. In contrast to performance techniques, reliability techniques seem to be added in a rather ad hoc fashion (e.g., most file systems apply sanity checks non-uniformly). Inclusion of a broader set of reliability techniques could harden all file systems.

Beyond these results, another outcome of our work is an annotated dataset of file-system patches, which we make publicly available for further study (at this URL: pages.cs.wisc.edu/~ll/fs-patch) by file-system developers, systems-language designers, and bug-finding tool builders. We show the utility of *PatchDB* by performing a case study (§6); specifically, we search the dataset to find bugs, performance fixes, and reliability techniques that are unusually common across all file systems. This example brings out one theme of our study, which is that there is a deep underlying similarity in Linux local file systems, even though these file systems are significantly different in nature (e.g., designs, features, and groups of developers). The commonalities we do find are good news: by studying past bug, performance, and reliability patches, and learning what issues and challenges lie therein, we can greatly improve the next generation of file systems and tools used to build them.

2 Methodology

In this section, we first give a brief description of our target file systems. Then, we illustrate how we analyze patches with a detailed example. Finally, we discuss the limitations of our methodology.

2.1 Target File Systems

Our goal in selecting a collection of disk-based file systems is to choose the most popular and important ones. The selected file systems should include diverse reliability features (e.g., physical journaling, logical journaling, checksumming, copy-on-write), data structures (e.g., hash tables, indirect blocks, extent maps, trees), performance optimizations (e.g., asynchronous thread pools, scalable algorithms, caching, block allocation for SSD devices), advanced features (e.g., pre-allocation, snapshot, resize, volumes), and even a range of maturity (e.g., stable, under development). For these reasons, we selected six file systems and their related modules: Ext3 with JBD [47], Ext4 with JBD2 [31], XFS [46], Btrfs [30], ReiserFS [13], and JFS [10]. Ext3, JFS, ReiserFS and XFS were all stable and in production use before the Linux 2.6 kernel. Ext4 was introduced in Linux 2.6.19 and marked stable in Linux 2.6.28. Btrfs was added into Linux 2.6.29 and is still under active development.

2.2 Classification of File System Patches

For each file system, we conduct a comprehensive study of its evolution by examining all patches from Linux 2.6.0 (Dec '03) to 2.6.39 (May '11). These are Linux main-line versions, which are released every three months with aggregate changes included in change logs. Patches consist of all formal modifications in each new kernel version, including new features, code maintenance, and bug fixes, and usually contain clear descriptions of their purpose and rich diagnostic information. On the other hand, Linux Bugzilla [3] and mailing lists [4, 5] are not as well organized as final patches, and may only contain a subset or superset of final changes merged in kernel.

To better understand the evolution of different file systems, we conduct a broad study to answer three categories of fundamental questions:

- *Overview*: What are the common types of patches in file systems and how do patches change as file systems evolve? Do patches of different types have different sizes?
- *Bugs*: What types of bugs appear in file systems? Do some components of file systems contain more bugs than others? What types of consequences do different bugs have?
- *Performance and Reliability*: What techniques are used by file systems to improve performance? What common reliability enhancements are proposed in file systems?

[PATCH] fix possible NULL pointer in fs/ext3/super.c.

In fs/ext3/super.c::ext3_get_journal() at line 1675
'journal' can be NULL, but it is not handled right
(detect by Coverity's checker).

```
--- /fs/ext3/super.c
+++ /fs/ext3/super.c
@@ -1675,6 +1675,7 @@ journal_t *ext3_get_journal()

1  if (!journal){
2      printk(KERN_ERR "EXT3: Could not load ... ");
3      iput(journal_inode);
4 +     return NULL;
5  }
6  journal->j_private = sb;
```

Figure 1: An Example Patch. An Ext3 patch.

To answer these questions, we manually analyzed each patch to understand its purpose and functionality, examining 5079 patches from the selected Linux 2.6 file systems. Each patch contains a patch header, a description body, and source-code changes. The patch header is a high-level summary of the functionality of the patch (e.g., fixing a bug). The body contains more detail, such as steps to reproduce the bug, system configuration information, proposed solutions, and so forth. Given these details and our knowledge of file systems, we categorize each patch along a number of different axes, as described later.

Figure 1 shows a real Ext3 patch. We can infer from the header that this patch fixes a null-pointer dereference bug. The body explains the cause of the null-pointer dereference and the location within the code. The patch also indicates that the bug was detected with Coverity [9].

This patch is classified as a bug (type=bug). The size is 1 (size=1) as one line of code is added. From the related source file (super.c), we infer the bug belongs to Ext3's superblock management (data-structure=super). A null-pointer access is a memory bug (pattern=memory,nullptr) and can lead to a crash (consequence=crash).

However, some patches have less information, making our analysis harder. In these cases, we sought out other sources of information, including design documents, forum and mailing-list discussions, and source-code analysis. Most patches are analyzed with high confidence given all the available information and our domain knowledge. Examples are shown throughout to give more insight as to how the classification is performed.

Limitations: Our study is limited by the file systems we chose, which may not reflect the characteristics of other file systems, such as other non-Linux file systems and flash-device file systems. We only examined kernel patches included in Linux 2.6 mainline versions, thus omitting patches for Ext3, JFS, ReiserFS, and XFS from Linux 2.4. As for bug representativeness, we only studied the bugs reported and fixed in patches, which is a biased subset; there may be (many) other bugs not yet reported. A similar study may be needed for user-space utilities, such as mkfs and fsck [33].

Type	Description
<i>Bug</i>	Fix existing bugs
<i>Performance</i>	Propose more efficient designs or implementations to improve performance (e.g., reducing synchronization overhead or use tree structures)
<i>Reliability</i>	Improve file-system robustness (e.g., data integrity verification, user/kernel pointer annotations, access-permission checking)
<i>Feature</i>	Implement new features
<i>Maintenance</i>	Maintain the code and documentation (e.g., adding documentation, fix compiling error, changing APIs)

Table 1: Patch Type. This table describes the classification and definition of file-system patches.

3 Patch Overview

File systems evolve through patches. A large number of patches are discussed and submitted to mailing lists, bug-report websites, and other forums. Some are used to implement new features, while others fix existing bugs. In this section, we investigate three general questions regarding file-system patches. First, what are file-system patch types? Second, how do patches change over time? Lastly, what is the distribution of patch sizes?

3.1 Patch Type

We classify patches into five categories (Table 1): bug fixes (*bug*), performance improvements (*performance*), reliability enhancements (*reliability*), new features (*feature*), and maintenance and refactoring (*maintenance*). Each patch usually belongs to a single category.

Figure 2(a) shows the number and relative percentages of patch types for each file system. Note that even though file systems exhibit significantly different levels of patch activity (shown by the total number of patches), the percentage breakdowns of patch types are relatively similar.

Maintenance patches are the largest group across all file systems (except Btrfs, a recent and not-yet-stable file system). These patches include changes to improve readability, simplify structure, and utilize cleaner abstractions; in general, these patches represent the necessary costs of keeping a complex open-source system well-maintained. Because maintenance patches are relatively uninteresting, we do not examine them further.

Bug patches have a significant presence, comprising nearly 40% of patches. Not surprisingly, the Btrfs has a larger percentage of bug patches than others; however, stable and mature file systems (such as Ext3) also have a sizable percentage of bug patches, indicating that bug fixing is a constant in a file system's lifetime (Figure 5). Because this class of patch is critical for developers and tool builders, we characterize them in detail later (§4).

Both *performance* and *reliability* patches occur as well, although with less frequency than maintenance and bug patches. They reveal a variety of techniques used by different file systems, motivating further study (§5).




Figure 2: **Patch Type and Bug Pattern.** This figure shows the distribution of patch types and bug patterns. The total number of patches is on top of each bar.

Finally, *feature* patches account for a small percentage of total patches; as we will see, most of *feature* patches contain more lines of code than other patches.

Summary: Nearly half of total patches are for code maintenance and documentation; a significant number of bugs exist in not only new file systems, but also stable file systems; all file systems make special efforts to improve their performance and reliability; feature patches account for a relatively small percentage of total patches.

3.2 Patch Trend

File systems change over time, integrating new features, fixing bugs, and enhancing reliability and performance. Does the percentage of different patch types increase or decrease with time?

We studied the changes in patches over time and found few changes (not shown). While the number of patches per version increased in general, the percentage of maintenance, bug, reliability, performance, and feature patches remained relatively stable. Although there were a few notable exceptions (e.g., Btrfs had a time where a large number of performance patches were added), the statistics shown in the previous section are relatively good summaries of the behavior at any given time. Perhaps most interestingly, bug patches do not decrease over time; living code bases constantly incorporate bug fixes (see §4).

Summary: The patch percentages are relatively stable over time; newer file systems (e.g., Btrfs) deviate occasionally; bug patches do not diminish despite stability.

3.3 Patch Size

Patch size is one approximate way to quantify the complexity of a patch, and is defined here as the sum of lines of added and deleted by a patch. Figure 3 displays the size distribution of bug, performance, reliability, and feature patches. Most *bug* patches are small; 50% are less than 10 lines of code. However, more complex file systems




Figure 3: **Patch Size.** This figure shows the size distribution for different patch types, in terms of lines of modifications.

Name	Description
<i>balloc</i>	Data block allocation and deallocation
<i>dir</i>	Directory management
<i>extent</i>	Contiguous physical blocks mapping
<i>file</i>	File read and write operations
<i>inode</i>	Inode-related metadata management
<i>trans</i>	Journaling or other transactional support
<i>super</i>	Superblock-related metadata management
<i>tree</i>	Generic tree structure procedures
<i>other</i>	Other supporting components (e.g., xattr, ioctl, resize)

Table 2: **Logical Components.** This table shows the classification and definition of file-system logical components.

tend to have larger bug patches (e.g., Btrfs and XFS) (not shown due to lack of space). Interestingly, feature patches are significantly larger than other patch types. Over 50% of these patches have more than 100 lines of code; 5% have over 1000 lines of code.

Summary: Bug patches are generally small; complicated file systems have larger bug patches; reliability and performance patches are medium-sized; feature patches are significantly larger than other patch types.

4 File System Bugs

In this section, we study file-system bugs in detail to understand their patterns and consequences comprehensively. First, we show the distribution of bugs in file-system logical components. Second, we describe our bug pattern classification, bug trends, and bug consequences. Finally, we analyze each type of bug with a more detailed classification and a number of real examples.

4.1 Correlation Between Code and Bugs

The code complexity of file systems is growing. FFS had only 1200 lines of code [32]; modern systems are notably larger, including Ext4 (29K LOC), Btrfs (47K LOC), and XFS (64K LOC). Several fundamental questions are germane: How is the code distributed among different logical components? Where are the bugs? Does each logical component have an equal degree of complexity?

File systems generally have similar logical components, such as inodes, superblocks, and journals. To en-




Figure 4: **File System Code and Bug Correlation.** This figure shows the correlation between code and bugs. The x-axis shows the average percent of code of each component (over all versions); the y-axis shows the percent of bugs of each component (over all versions).

able fair comparison, we partition each file system into nine logical components (Table 2).

Figure 4 shows the percentage of bugs versus the percentage of code for each of the logical components across all file systems and versions. Within a plot, if a point is above the $y = x$ line, it means that a logical component (e.g., inodes) has more than its expected share of bugs, hinting at its complexity; a point below said line indicates a component (e.g., a tree) with relatively few bugs per line of code, thus hinting at its relative ease of implementation.

We make the following observations. First, for all file systems, the *file*, *inode*, and *super* components have a high bug density. The file component is high in bug density either due to bugs on the fsync path (Ext3) or custom file I/O routines added for higher performance (XFS, Ext4, ReiserFS, JFS), particularly so for XFS, which has a custom buffer cache and I/O manager for scalability [46]. The inode and superblock are core metadata structures with rich and important information for files and file systems, which are widely accessed and updated; thus, it is perhaps unsurprising that a large number of bugs arise therein (e.g., forgetting to update a time field in an inode, or not properly using a superblock configuration flag).

Type	Sub-Type	Description
Semantic	State	Incorrectly update or check file-system state
	Logic	Wrong algorithm/assumption/implementation
	Config	Missed configuration
	I/O Timing	Wrong I/O requests order
Concurrency	Generic	Generic semantic bugs: wrong type, typo
	Atomicity	The atomic property for accesses is violated
	Order	The order of multiple accesses is violated
	Deadlock	Deadlock due to wrong locking order
	Miss unlock	Miss a paired unlock
	Double unlock	Unlock twice
	Wrong lock	Use the wrong lock
Memory	Resource leak	Fail to release memory resource
	Null pointer	Dereference null pointer
	Dangling Pt	Dereference freed memory
	Uninit read	Read uninitialized variables
	Double free	Free memory pointer twice
	Buf overflow	Overrun a buffer boundary
Error Code	Miss Error	Error code is not returned or checked
	Wrong Error	Return or check wrong error code

Table 3: **Bug Pattern Classification.** This table shows the classification and definition of file-system bugs.

Second, transactional code represents a substantial percentage of each code base (as shown by the relatively high x-axis values) and, for most file systems, has a proportional amount of bugs. This relationship holds for Ext3 as well, even though Ext3 uses a separate journaling module (JBD); Ext4 (with JBD2) has a slightly lower percentage of bugs because it was built upon a more stable JBD from Linux 2.6.19. In summary, transactions continue to be a double-edged sword in file systems: while transactions improve data consistency in the presence of crashes, they often add many bugs due to their large code bases.

Third, the percentage of bugs in *tree* components of XFS, Btrfs, ReiserFS, and JFS is surprisingly small compared to code size. One reason may be the care taken to implement such trees (e.g., the tree code is the only portion of ReiserFS filled with assertions). File systems should be encouraged to use appropriate data structures, even if they are complex, because they do not induce an inordinate amount of bugs.

Although bug patches also relate to feature patches, it is difficult to correlate them precisely. Code changes partly or totally overlap each other overtime. A bug patch may involve both old code and recent feature patches.

Summary: The file, inode, and superblock components contain a disproportionately large number of bugs; transactional code is large and has a proportionate number of bugs; tree structures are not particularly error-prone, and should be used when needed without much worry.

4.2 Bug Patterns

To build a more reliable file system, it is important to understand the type of bugs that are most prevalent and the typical patterns across file systems. Since different types of bugs require different approaches to detect and




Figure 5: **Bug Pattern Evolution.** This figure shows the bug pattern evolution for each file system over all versions.

fix, these fine-grained bug patterns provide useful information to developers and tool builders alike.

We partition file-system bugs into four categories based on their root causes as shown in Table 3. The four major categories are *semantic* [26, 44], *concurrency* [16, 28], *memory* [14, 26, 44], and *error code* bugs [19, 40].

Figure 2(b) (page 4) shows the total number and percentage of each type of bug across file systems. There are about 1800 total bugs, providing a great opportunity to explore bug patterns at scale. Semantic bugs dominate other types (except for ReiserFS). Most semantic bugs require file-system domain knowledge to understand, detect, and fix; generic bug-finding tools (e.g., Coverity [9]) may have a hard time finding these bugs. Concurrency bugs account for about 20% on average across file systems (except for ReiserFS), providing a stark contrast to user-level software where fewer than 3% of bugs are concurrency-related [26, 42, 45]. ReiserFS stands out along these measures because of its transition, in Linux 2.6.33, away from the Big Kernel Lock (BKL), which introduced a large number of concurrency bugs. There are also a fair number of memory-related bugs in all file systems; their percentages are lower than that reported in user-level software [26, 45]. Many research and commercial tools have

Type	Description
<i>Corruption</i>	On-disk or in-memory data structures are corrupted (e.g., file data or metadata corruption, wrong statistics)
<i>Crash</i>	File system becomes unusable (e.g., dereference null pointer, assertion failures, panics)
<i>Error</i>	Operation failure or unexpected error code returned (e.g., failed write operation due to ENOSPC error)
<i>Deadlock</i>	Wait for resources in circular chain
<i>Hang</i>	File system makes no progress (e.g., infinite loop, live lock)
<i>Leak</i>	System resources are not freed after usage (e.g., forget to free allocated file-system objects)
<i>Wrong</i>	Diverts from expectation, excluding the above ones (e.g., undefined behavior, security vulnerability)

Table 4: **Bug Consequence Classification.** This table shows the definitions of various bug consequences.

been developed to detect memory bugs [9, 35], and some of them are used to detect file-system bugs. Error code bugs account for only 10% of total bugs.

Summary: Beyond maintenance, bug fixes are the most common patch type; over half of file-system bugs are semantic bugs, likely requiring domain knowledge to find and fix; file systems have a higher percentage of concurrency bugs compared with user-level software; memory and error code bugs arise but in smaller percentages.

4.3 Bug Trends

File systems mature from the initial development stage to the stable stage over time, by applying bug-fixing, performance and reliability patches. Various bug detection and testing tools are also proposed to improve file-system stability. A natural question arises: do file-system bug patterns change over time, and in what way?

Our results (Figure 5) show that within bugs, the relative percentage of semantic, concurrency, memory, and error code bugs varies over time, but does not converge; a great example is XFS, which under constant development goes through various cycles of higher and lower numbers of bugs. Interesting exceptions occasionally arise (e.g., the BKL removal from ReiserFS led to a large increase in concurrency bugs in 2.6.33). JFS does experience a decline in bug patches, perhaps due to its decreasing usage and development [6]. JFS and ReiserFS both have relatively small developer and user bases compared to the more active file systems XFS, Ext4 and Btrfs.

Summary: Bug patterns do not change significantly over time, increasing and decreasing cyclically; large deviations arise due to major structural changes.

4.4 Bug Consequences

As shown in Figure 2(b) (on page 4), there are a significant number of bugs in file systems. But how serious are these file-system bugs? We now categorize each bug by impact; such *bug consequences* include severe ones (data corruption, system crashes, unexpected errors, deadlocks,




Figure 6: Bug Consequences. This figure displays the breakdown of bug consequences for file systems and bug patterns. The total number of consequences is shown on top of each bar. A single bug may cause multiple consequences; thus, the number of consequences instances is slightly higher than that of bugs in Figure 2(b).

system hangs and resource leaks), and other wrong behaviors. Table 4 provides more detail on these categories.

Figure 6(a) shows the per-system breakdowns. Data corruption is the most predominant consequence (40%), even for well-tested and mature file systems. Crashes account for the second largest percentage (20%); most crashes are caused by explicit calls to `BUG()` or `Assert()` as well as null-pointer dereferences. If the patch mentions that the crash also causes corruption, then we classify this bug with multiple consequences. Unexpected errors and deadlocks occur quite frequently (just under 10% each on average), whereas other bug consequences arise less often. For example, exhibiting the wrong behavior without more serious consequences accounts for only 5-10% of consequences in file systems, whereas it is dominant in user applications [26].

Given that file-system bugs are serious bugs, we were curious: do certain bug types (e.g., semantic, concurrency, memory, or error code) exhibit different levels of severity? Figure 6(b) shows the relationship between consequences and bug patterns. Semantic bugs lead to a large percentage of corruptions, crashes, errors, hangs, and wrong behaviors. Concurrency bugs are responsible for nearly all deadlocks (almost by definition) and a fair percentage of corruptions and hangs. Memory bugs lead to many memory leaks (as expected) and a fair amount of crashes. Finally, error code bugs lead to a relatively small percentage of corruptions, crashes, and (unsurprisingly) errors.

Summary: File system bugs cause severe consequences; corruptions and crashes are most common; wrong behavior is uncommon; semantic bugs can lead to significant amounts of corruptions, crashes, errors, and hangs; all bug types have severe consequences.

4.5 Bug Pattern Examples and Analysis

To gain further insight into the different classes of bugs, we now describe each class in more detail. We present examples of each and further break down each major class (e.g., memory bugs) into smaller sub-classes (e.g., leaks, null-pointer dereferences, dangling pointers, uninitialized reads, double frees, and buffer overflows).

4.5.1 Semantic Bugs

Semantic bugs are dominant in file systems, as shown in Figure 2(b). Understanding these bugs often requires file-system domain knowledge. Semantic bugs usually are difficult to categorize in an informative and general way. However, we are the first to identify several common types of file-system specific semantic bugs based on extensive analysis and careful generalization of many semantic bugs across file systems. These common types and typical patterns provide useful guidelines for analysis and detection of file-system semantic bugs. We partition the semantic bugs into five categories as described in Table 3, including *state*, *logic*, *config*, *I/O timing* and *generic*. Figure 7(a) shows the percentage breakdown and total number of semantic bugs; each is explained in detail below.

File systems maintain a large amount of in-memory and on-disk state. Generally, operations transform the file system from one consistent state to another; a mistaken state update or access may lead to serious consequences. As shown in Figure 7(a), these *state* bugs contribute to roughly 40% of semantic bugs.

An example of a *state* bug is shown in S1 of Table 5 (on page 9), which misses an inode-field update. Specifically, the buggy version of `ext3_rename()` does not update the `mtime` and `ctime` of the directory into which the file is moved, leaving metadata in an incorrect state.

There are also numerous *logic* bugs, which arise via the use of wrong algorithms, bad assumptions, and incorrect implementations. An example of a wrong algorithm is shown in S2 of Table 5: `find_group_other()` tries to find a block group for inode allocation, but does not check all candidate groups; the result is a possible `ENOSPC` error even when the file system has free inodes.

File system behavior is also affected by various configuration parameters, such as mount options and special hardware support. Unfortunately, file systems often forget or misuse such configuration information (about 10% to 15% of semantic bugs are of this flavor). A semantic configuration bug is shown in S3 of Table 5; when Ext4 loads the journal from disk, it forgets to check if the device is read-only before updating the on-disk superblock.

Correct I/O request ordering is critical for crash consistency in file systems. The *I/O timing* category contains bugs involving incorrect I/O ordering. For example, in ordered journal mode, a bug may flush metadata to disk before the related data blocks are persisted. We found




Figure 7: **Detailed Bug Patterns.** The detailed classification for each bug pattern; total number of bugs is shown on top of each bar.

that only a small percentage of semantic bugs (3-9%) are I/O timing bugs; however, these bugs can lead to potential data loss or corruption.

A fair amount of *generic* bugs also exist in all file systems, such as using the wrong variable type or simple typos. These bugs are general coding mistakes (such as comparing unsigned variable with zero [48]), and may be fixed without much file-system knowledge.

Summary: Incorrect state update and logic mistakes dominate semantic bug patterns; configuration errors are also not uncommon; incorrect I/O orderings are rare (but can have serious consequences); generic bugs require the least file-system knowledge to understand.

4.5.2 Concurrency Bugs

Concurrency bugs have attracted a fair amount of attention in the research community as of late [16, 22, 28, 49, 50]. To better understand file-system concurrency bugs, we classify them into six types as shown in Table 3 (on page 5): *atomicity violations*, *deadlocks*, *order violations*, *missed unlocks*, *double unlocks*, and *wrong locks*.

Figure 7(b) shows the percentage and total number of each category of concurrency bugs. Atomicity violation bugs are usually caused by a lack of proper synchronization methods to ensure exclusive data access, often leading to data corruption.

An example of an atomicity violation bug in Ext4 is shown in C1 of Table 5. For this bug, when two CPUs simultaneously allocate blocks, there is no protection for the `i_cached_extent` structure; this atomicity violation could thus cause the wrong location on disk to be read or written. A simple spin-lock resolves the bug.

There are a large number of *deadlocks* in file systems (about 40%). Two typical causes are the use of the wrong kernel memory allocation flag and calling a blocking function when holding a spin lock. These patterns are not common in application-level deadlocks, and thus are useful to both developers (who should be wary of such patterns) and tool builders (who should detect them).

Many deadlocks are found in ReiserFS, once again due to the BKL. The BKL could be acquired recursively; replacing it introduced a multitude of locking violations, many of which led to deadlock.

A typical memory-related deadlock is shown in C2 of Table 5. Btrfs uses `extent_readpages()` to read free space information; however, it should not use `GFP_KERNEL` flag to allocate memory, since the VM memory allocator `kswapd` will recursively call into file-system code to free memory. The fix changes the flag to `GFP_NOFS` to prevent VM re-entry into file-system code.

The remaining four categories account for a small percentage. Missing unlocks happen mostly in exit or failure paths (e.g., putting resource releases at the end of functions with `goto` statements). C3 of Table 5 shows a missing-unlock bug. `ext3_group_add()` locks super block (line 1) but forgets to unlock on an error (line 4).

Summary: Concurrency bugs are much more common in file systems than in user-level software. Atomicity and deadlock bugs represent a significant majority of concurrency bugs; many deadlock bugs are caused by wrong kernel memory-allocation flags; most missing unlocks happen on exit or failure paths.

4.5.3 Memory Bugs

Memory-related bugs are common in many source bases, and not surprisingly have been the focus of many bug detection tools [9, 35]. We classify memory bugs into six categories, as shown in Table 3: *resource leaks*, *null pointer dereferences*, *dangling pointers*, *uninitialized reads*, *double frees*, and *buffer overflows*.

Resource leaks are the most dominant, over 40% in aggregate; in contrast, studies of user-level programs show notably lower percentages [26, 42, 45]. We find that roughly 70% of resource leaks happen on exit or failure paths; we investigate this further later (§4.6).

An example of resource leaks (M1 of Table 5) is found in `btrfs_new_inode()` which allocates an inode but forgets to free it upon failure.

<i>ext3/namei.c, 2.6.26</i> 1 ext3.rename(...){ 2 + new_dir->i_ctime = CURRENT.TIME_SEC; 3 + new_dir->i_mtime = CURRENT.TIME_SEC; 4 + ext3.mark_inode_dirty(handle, new_dir);	Semantic (S1)	<i>ext3/lalloc.c, 2.6.4</i> 1 find_group_other(...){ 2 - group = parent.group + 1; 3 - for (i = 2; i < ngroups; i++) { 4 + group = parent.group; 5 + for (i = 0; i < ngroups; i++) {	Semantic (S2)
<i>ext4/super.c, 2.6.37</i> 1 ext4.load_journal(...){ 2 - if (journal.devnum && ...) 3 + if (!read_only && journal.devnum ...) 4 es->s_journal_dev = devnum;	Semantic (S3)	<i>ext4/extents.c, 2.6.30</i> 1 ext4_ext_put_in_cache(...){ 2 + spin_lock(i_block_reservation_lock); 3 cex = &EXT4_I(inode)->i_cached_extent; 4...6 cex->ec_FOO = FOO; // elided for brevity 7 + spin_unlock(i_block_reservation_lock);	Concurrency (C1)
<i>btrfs/extent_io.c, 2.6.39</i> 1 extent_readpages(...){ 2 if (!add_to_page_cache_lru(page, mapping, 3 - page->index, GFP_KERNEL)) { 4 + page->index, GFP_NOFS)) { 5 _extent_read_full_page(...);	Concurrency (C2)	<i>ext3/resize.c, 2.6.17</i> 1 lock_super(sb); 2 if (input->group != sbi->s_groups_count){ 3 4 + unlock_super(sb); 5 err = -EBUSY; 6 goto exit_journal;	Concurrency (C3)
<i>btrfs/inode.c, 2.6.30</i> 1 btrfs_new_inode(...){ 2 inode = new_inode(...); 3 ret = btrfs_set_inode_index(...); 4 - if (ret) 5 - return ERR_PTR(ret); 6 + if (ret) { 7 + iput(inode); return ERR_PTR(ret); 8 + }	Memory (M1)	<i>ext3/super.c, 2.6.7</i> 1 ext3_get_journal(...){ 2 if (!journal) { 3 4 + return NULL; 5 } 6 journal->j_private = sb;	Memory (M2)
<i>reiserfs/xattr_acl.c, 2.6.16</i> 1 reiserfs_get_acl(...){ 2 acl = posix_acl_from_disk(...); 3 - *p_acl = posix_acl_dup(acl); 4 + if (!IS_ERR(acl)) 5 + *p_acl = posix_acl_dup(acl);	Error Code (E1)	<i>jfs/jfs_imap.c, 2.6.27</i> 1 diAlloc(...){ 2 jfs_error(...); 3 - return EIO; 4 + return -EIO;	Error Code (E2)
<i>ext4/extents.c, 2.6.31</i> 1 ext4_fiemap(...){ 2 - down_write(&EXT4_I(inode)->i_data_sem); 3 + down_read(&EXT4_I(inode)->i_data_sem); 4 error = ext4_ext_walk_space(...); 5 - up_write(&EXT4_I(inode)->i_data_sem); 6 + up_read(&EXT4_I(inode)->i_data_sem);	Performance (P1)	<i>btrfs/free-space-cache.c, 2.6.39</i> 1 btrfs_find_space_cluster(...){ 2 + if (bg->free_space < minbytes){ 3 + spin_unlock(&bg->tree_lock); 4 + return -ENOSPC; 5 + } 6 /* start to search for blocks */	Performance (P2)

Table 5: **Code Examples.** This table shows the code examples of bug patterns and performance patches.

As we see in Figure 7(c), null-pointer dereferences are also common in both mature and young file systems (the remaining memory bugs account for small percentages). An example is shown in M2 of Table 5; a return statement is missing, leading to a null-pointer dereference.

Summary: Resource leaks are the largest category of memory bug, significantly higher than that in user-level applications; null-pointer dereferences are also common; failure paths contribute strongly to these bugs; many of these bugs have simple fixes.

4.5.4 Error Code Bugs

File systems need to handle a wide range of errors, including memory-allocation failures, disk-block allocation failures, I/O failures [7, 8], and silent data corruption [37]. Handling such faults, and passing error codes through a complex code base, has proven challenging [19, 40]. Here, we further break down error-code errors.

We partition the error code bugs into *missing error codes* and *wrong error codes* as described in Table 3. Figure 7(d) shows the breakdown of error code bugs. Missing errors are generally twice as prevalent as wrong errors (except for JFS, which has few of these bugs overall).

A missing error code example is shown in E1 of Table 5. The routine `posix_acl_from_disk()` could re-

turn an error code (line 2). However, without error checking, `acl` is accessed and thus the kernel crashes (line 3).

An example of a wrong error code is shown in E2 of Table 5. `diAlloc()`'s return value should be `-EIO`. However, in line 3, the original code returns the close (but wrong) error code `EIO`; callers thus fail to detect the error.

Summary: Error handling bugs occur in two flavors, missing error handling or incorrect error handling; the bugs are relatively simple in nature.

4.6 The Failure Path

Many bugs we found arose not in common-case code paths but rather in more unusual fault-handling cases [19, 52]. This type of error handling (i.e., reacting to disk or memory failures) is critical to robustness, since bugs on failure paths can lead to serious consequences. We now quantify bug occurrences on failure paths; Tables 6 (a) and (b) present our accumulated results.

As we can see from the first table, roughly a third of bugs are introduced on failure paths across all file systems. Even mature file systems such as Ext3 and XFS make a significant number of mistakes on these rarer code paths.

When broken down by bug type in the second table, we see that roughly a quarter of semantic bugs occur on failure paths, usually in the previously-defined *state* and *logic*

XFS	Ext4	Btrfs	Ext3	ReiserFS	JFS
200 (39.1%)	149 (33.1%)	144 (40.2%)	88 (38.4%)	63 (39.9%)	28 (35%)
(a) By File System					
Semantic	Concurrency	Memory	Error Code		
283 (27.7%)	93 (25.4%)	117 (53.4%)	179 (100%)		
(b) By Bug Pattern					

Table 6: **Failure Related Bugs.** This table shows the number and percentage of the bugs related to failures in file systems.

categories. Once a failure happens (e.g., an I/O fails), the file system needs to free allocated disk resources and update related metadata properly; however, it is easy to forget these updates, or perform them incorrectly, leading to many *state* bugs. In addition, wrong algorithms (*logic* bugs) are common; for example, when block allocation fails, most file systems return ENOSPC immediately instead of retrying after committing buffered transactions.

A quarter of concurrency bugs arise on failure paths. Sometimes, file systems forget to unlock locks, resulting in deadlock. Moreover, when file systems output errors to users, they sometimes forget to unlock before calling blocking error-output functions (deadlock). These types of mistakes rarely arise in user-level code [28].

For memory bugs, most resource-leak bugs stem from forgetting to release allocated resources when I/O or other failures happen. There are also numerous null-pointer dereference bugs which incorrectly assume certain pointers are still valid after a failure. Finally (and obviously), all error code bugs occur on failure paths (by definition).

It is difficult to fully test failure-handling paths to find all types of bugs. Most previous work has focused on memory resource leaks [41, 52], missing unlock [41, 52] and error codes [19, 40]; however, existing work can only detect a small portion of failure-handling errors, especially omitting a large amount of semantic bugs on failure paths. Our results provide strong motivation for improving the quality of failure-handling code in file systems.

Summary: A high fraction of bugs occur due to improper behavior in the presence of failures or errors across all file systems; memory-related errors are particularly common along these rarely-executed code paths; a quarter of semantic bugs are found on failure paths.

5 Performance and Reliability

A small but important set of patches improve performance and reliability, which are quantitatively different than bug patches (Figure 3). Performance and reliability patches account for 8% and 7% of patches respectively.

5.1 Performance Patches

Performance is critical for all file systems. Performance patches are proposed to improve existing designs or implementations. We partition these patches into six categories as shown in Table 7, including synchronization (*sync*), access optimization (*access*), scheduling (*sched*),

Type	Description
<i>Synchronization</i>	Inefficient usage of synchronization methods (e.g., removing unnecessary locks, using smaller locks, using read/write locks)
<i>Access Optimization</i>	Apply smarter access strategies (e.g., caching metadata and statistics, avoiding unnecessary I/O and computing)
<i>Schedule</i>	Improve I/O operations scheduling (e.g., batching writes, opportunistic readahead)
<i>Scalability</i>	Scale on-disk and in-memory data structures (e.g., using trees or hash tables, per block group structures, reducing memory usage of inodes)
<i>Locality</i>	Overcome sub-optimal data block allocations (e.g., reducing file fragmentation, clustered I/Os)
<i>Other</i>	Other performance improvement techniques (e.g., reducing function stack usage)

Table 7: **Performance Patch Type.** This table shows the classification and definition of performance patches.

scalability (*scale*), locality (*locality*), and *other*. Figure 8(a) shows the breakdown.

Synchronization-based performance improvements account for over a quarter of all performance patches across file systems. Typical solutions used include removing a pair of unnecessary locks, using finer-grained locking, and replacing write locks with read/write locks. A *sync* patch is shown in P1 of Table 5; `ext4_fiemap()` uses write instead of read semaphores, limiting concurrency.

Access patches use smarter strategies to optimize performance, including caching and work avoidance. For example, Ext3 caches metadata stats in memory, avoiding I/O. Figure 8(a) shows *access* patches are popular. An example Btrfs *access* patch is shown in P2 of Table 5; before searching for free blocks, it first checks whether there is enough free space, avoiding unnecessary work.

Sched patches improve I/O scheduling for better performance, such as batching of writes, opportunistic readahead, and avoiding unnecessary synchrony in I/O. As can be seen, *sched* has a similar percentage compared to *sync* and *access*. *Scale* patches utilize scalable on-disk and in-memory data structures, such as hash tables, trees, and per block-group structures. XFS has a large number of *scale* patches, as scalability was always its priority.

Summary: Performance patches exist in all file systems; *sync*, *access*, and *sched* each account for a quarter of the total; many of the techniques used are fairly standard (e.g., removing locks); while studying new synchronization primitives, we should not forget about performance.

5.2 Reliability Patches

Finally we study our last class of patch, those that aim to improve file-system reliability. Different from bug-fix patches, reliability patches are not utilized for correctness. Rather, for example, such a patch may check whether the super block is corrupted before mounting the file system; further, a reliability patch might enhance error propagation [19] or add more debugging information. Table 8 presents the classification of these *reliability* patches, including adding assertions and other functional robustness

Type	Description
<i>Robust</i>	Enhance file-system robustness (e.g., boundary limits and access permission checking, additional internal assertions)
<i>Corruption Defense</i>	Improve file systems’ ability to handle various possible corruptions
<i>Error Enhancement</i>	Improve original error handling (e.g., gracefully handling failures, more detailed error codes)
<i>Annotation</i>	Add endianness, user/kernel space pointer and lock annotations for early bug detection
<i>Debug</i>	Add more internal debugging or tracing support

Table 8: **Reliability Patch Type.** This table shows the classification and definition of reliability patches.

(*robust*), corruption defense (*corruption*), error enhancement (*error*), annotation (*annotation*), and debugging (*debug*). Figure 8(b) displays the distributions.

Robust patches check permissions, enforce file-system limits, and handle extreme cases in a more friendly manner. Btrfs has the largest percentage of these patches, likely due to its early stage of development.

Corruption defense patches validate the integrity of metadata when reading from disk. For example, a patch to the JBD (used by Ext3) checks that the journal length is valid before performing recovery; similarly, a patch to Ext4 checks that a directory entry is valid before traversing that directory. In general, many *corruption* patches are found at the I/O boundary, when reading from disk.

Error enhancement patches improve error handling in a variety of ways, such as more detail in error codes, removing unnecessary error messages, and improving availability, for example by remounting read-only instead of crashing. This last class is common in all file systems, which each slowly replaced unnecessary `BUG()` and assertion statements with more graceful error handling.

Annotation patches label variables with additional type information (e.g., endianness) and locking rules to enable better static checking. ReiserFS uses lock annotations to help prevent deadlock, whereas XFS uses endianness annotations for numerous variable types. *Debug* patches simply add more diagnostic information at failure-handling points within the file system.

Interestingly, reliability patches appear more ad hoc than bug patches. For bug patches, most file systems have similar pattern breakdowns. In contrast, file systems make different choices for reliability, and do so in a generally non-uniform manner. For example, Btrfs focuses more on *Robust* patches, while Ext3 and Ext4 prefer to add more *Corruption* defense patches.

Summary: We find that *reliability* patches are added to file systems over time as part of hardening; most add simple checks, defend against corruption upon reading from disk, or improve availability by returning errors instead of crashing; annotations help find problems at compile time; debug patches add diagnostic information; reliability patch usage, across all file systems, seems ad hoc.




Figure 8: **Performance and Reliability Patches.** This figure shows the performance and reliability patterns. The total number of patches is shown on top of each bar.

6 Case Study Using PatchDB

The patch dataset constructed from our analysis of 5079 patches contains fine-grained information, including characterization of bug patterns (e.g., which semantic bugs forget to synchronize data), detailed bug consequences (e.g., crashes caused by assertion failures or null-pointer dereferences), incorrect bug fixes (e.g., patches that are reverted after being accepted), performance techniques (e.g., how many performance patches remove unnecessary locks), and reliability enhancements (e.g., the location of metadata integrity checks). These details enable further study to improve file-system designs, propose new system language constructs, build custom bug-detection tools, and perform realistic fault injection.

In this section, we show the utility of *PatchDB* by examining which patches are common across all file systems. Due to space concerns, we only highlight a few interesting cases. A summary is found in Table 9.

We first discuss specific common bugs. Within semantic bugs is *forget sync*, in which a file system forgets to force data or metadata to disk. Most *forget sync* bugs relate to *fsync*. Even for stable file systems, there are a noticeable number of these bugs, leading to data loss or corruption under power failures. Another common mistake is *forget config*, in which mount options, feature sets, or hardware support are overlooked. File systems also return the `ENOSPC` error code despite the presence of free blocks (*early enospc*); Btrfs has the largest number of these bugs, and even refers to the Ext3 fix strategy in its patches. Even though semantic bugs are dominant in file systems, few tools can detect semantic bugs due to the difficulty of specifying correct behavior [15, 25, 27]. Fortunately, we find that many semantic bugs appear across file systems, which can be leveraged to improve bug detection.

For concurrency bugs, forgetting to lock an inode when updating it is common; perhaps a form of monitors [20]

Patch Type	Typical Cases	XFS	Ext4	Btrfs	Ext3	Reiser	JFS
Semantic	forget sync	17	11	6	11	5	1
	forget config	43	43	23	16	8	1
	early enospc	5	9	14	7		
	wrong log credit	6	8	1	1	1	
Concurrency	lock inode update	6	5	2	4	4	2
	lock sleep	8	8	1	1	8	
	wrong kmalloc flag	20	3	3	2		1
	miss unlock	10	7	4	2	2	4
Memory	leak on failure	14	21	16	11	1	3
	leak on exit	1		1	4		1
Error Code	miss I/O error	10	11	8	15	4	1
	miss mem error	4	2	13	1	1	
	bad error access		3	8			2
Performance	remove lock	17	14	14	8	5	1
	avoid redun write	6	4	5	4		2
	check before work	8	5	15	2	1	
	save struct mem	3	9		1	3	
Reliability	metadata validation	12	9	1	7	2	1
	graceful handle	8	6	5	5	1	4

Table 9: **Common File System Patches.** *This table shows the classification and count of common patches across all file systems.*

would help. Calling a blocking function when holding a spin lock (*lock sleep*) occurs frequently (also in drivers [14, 36]). As we saw earlier (§4.5.2), using the wrong kernel memory allocation flag is a major source of deadlock (particularly XFS). All file systems miss unlocks frequently, in contrast to user applications [28].

For memory bugs, leaks happen on failure or exit paths frequently. For error code bugs, there are a large number of *missed I/O error* bugs. For example, Ext3, JFS, ReiserFS and XFS all ignore write I/O errors on fsync before Linux 2.6.9 [37]; as a result, data could be lost even when fsync returned successfully. Memory allocation errors are also often ignored (especially in Btrfs). Three file systems mistakenly dereference error codes.

For performance patches, removing locks (without sacrificing correctness) is common. File systems also tend to write redundant data (e.g., *fdatasync* unnecessarily flushes metadata). Another common performance improvement case is *check before work*, in which missing specific condition checking costs unnecessary I/O or CPU overhead.

Finally, for reliability patches, metadata validation (i.e., inode, super block, directory and journal) is popular. Most of these patches occur in similar places (e.g., when mounting the file system, recovering from the journal, or reading an inode). Also common is replacing `BUG()` and `Assert()` calls with more graceful error handling.

Summary: Despite their diversity, file-system patches share many similarities across implementations; some examples occur quite frequently; PatchDB affords new opportunities to study such phenomena in great detail.

7 Related Work

Operating-System Bugs: Faults in Linux have been studied [14, 36]. Static analysis tools are used to find poten-

tial bugs in Linux 1.0 to 2.4.1 [14] and Linux 2.6.0 to 2.6.33 [36]. Most detected faults are generic memory and concurrency bugs. Both studies find that device drivers contain the most faults, while Palix et al. [36] also show that file-system errors are rising. Yin et al. [53] analyze incorrect bug-fixes in several operating systems. Our work embellishes these studies, focusing on all file-system bugs found and fixed over eight years and providing more detail on which bugs plague file systems.

User-Level Bugs: Various aspects of modern user-level open source software bugs have also been studied, including patterns, impacts, reproducibility, and fixes [16, 26, 28, 42, 50]. As our findings show, file-systems bugs display different characteristics compared with user-level software bugs, both in their patterns and consequences (e.g., file-system bugs have more serious consequences than user-level bugs; concurrency bugs are much more common). One other major difference is scale; the number of bugs (about 1800) we study is larger than previous efforts [16, 26, 28, 42, 50].

File-System Bugs: Several research projects have been proposed to detect and analyze file-system bugs. For example, Yang et al. [51, 52] use model checking to detect file-system errors; Gunawi et al. [19] use static analysis techniques to determine how error codes are propagated in file systems; Rubio-Gonzalez et al. [40] utilize static analysis to detect similar problems; Prabhakaran et al. [37] study how file systems handle injected failures and corruptions. Our work complements this work with insights on bug patterns and root causes. Further, our public bug dataset provides useful hints and patterns to aid in the development of new file-system bug-detection tools.

8 Conclusions

We performed a comprehensive study of 5079 patches across six Linux file systems; our analysis includes one of the largest studies of bugs to date (nearly 1800 bugs). Our observations, summarized in the introduction and throughout, should be of utility to file-system developers, systems-language designers, and tool makers; the careful study of these results should result in a new generation of more robust, reliable, and performant file systems.

Acknowledgments

We thank Ric Wheeler (our shepherd) and the anonymous reviewers for their excellent feedback and suggestions. We also thank the members of the ADSL research group for their insightful comments. This material is based upon work supported by the National Science Foundation under the following grants: CNS-1218405, CCF-0937959, CSR-1017518, CCF-1016924, as well as generous support from NetApp, EMC, and Google. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF or other institutions.

References

- [1] Coverity Scan: 2011 Open Source Integrity Report. <http://www.coverity.com/library/pdf/coverity-scan-2011-open-source-integrity-report.pdf>.
- [2] First Galaxy Nexus Rom Available, Features Ext4 Support. <http://androidspin.com/2011/12/06/first-galaxy-nexus-rom-available-features-ext4-support/>.
- [3] Kernel Bug Tracker. <http://bugzilla.kernel.org/>.
- [4] Linux Filesystem Development List. <http://marc.info/?l=linux-fsdevel>.
- [5] Linux Kernel Mailing List. <http://lkml.org/>.
- [6] IBM Journaled File System. [http://en.wikipedia.org/wiki/JFS_\(file_system\)](http://en.wikipedia.org/wiki/JFS_(file_system)), September 2012.
- [7] Lakshmi N. Bairavasundaram, Garth R. Goodson, Shankar Pasupathy, and Jiri Schindler. An Analysis of Latent Sector Errors in Disk Drives. In *Proceedings of the 2007 ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '07)*, San Diego, California, June 2007.
- [8] Lakshmi N. Bairavasundaram, Garth R. Goodson, Bianca Schroeder, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. An Analysis of Data Corruption in the Storage Stack. In *Proceedings of the 6th USENIX Symposium on File and Storage Technologies (FAST '08)*, pages 223–238, San Jose, California, February 2008.
- [9] Al Bessey, Ken Block, Ben Chelf, Andy Chou, Bryan Fulton, Seth Hallem, Charles Henri-Gros, Asya Kamsky, Scott McPeak, and Dawson Engler. A Few Billion Lines of Code Later: Using Static Analysis to Find Bugs in the Real World. *Communications of the ACM*, February 2010.
- [10] Steve Best. JFS Overview. <http://jfs.sourceforge.net/project/pub/jfs.pdf>, 2000.
- [11] Simona Boboila and Peter Desnoyers. Write Endurance in Flash Drives: Measurements and Analysis. In *Proceedings of the 8th USENIX Symposium on File and Storage Technologies (FAST '10)*, San Jose, California, February 2010.
- [12] Jeff Bonwick and Bill Moore. ZFS: The Last Word in File Systems. http://opensolaris.org/os/community/zfs/docs/zfs_last.pdf, 2007.
- [13] Florian Buchholz. The structure of the Reiser file system. <http://homes.cerias.purdue.edu/~florian/reiser/reiserfs.php>, January 2006.
- [14] Andy Chou, Junfeng Yang, Benjamin Chelf, Seth Hallem, and Dawson Engler. An Empirical Study of Operating System Errors. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01)*, pages 73–88, Banff, Canada, October 2001.
- [15] Dawson Engler, David Yu Chen, Seth Hallem, Andy Chou, and Benjamin Chelf. Bugs as Deviant Behavior: A General Approach to Inferring Errors in Systems Code. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01)*, pages 57–72, Banff, Canada, October 2001.
- [16] Pedro Fonseca, Cheng Li, Vishal Singhal, and Rodrigo Rodrigues. A Study of the Internal and External Effects of Concurrency Bugs. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN '10)*, Chicago, USA, June 2010.
- [17] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The Google File System. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP '03)*, pages 29–43, Bolton Landing, New York, October 2003.
- [18] L. M. Grupp, A. M. Caulfield, J. Coburn, S. Swanson, E. Yaakobi, P. H. Siegel, and J. K. Wolf. Characterizing Flash Memory: Anomalies, Observations, and Applications. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'09)*, New York, New York, December 2009.
- [19] Haryadi S. Gunawi, Cindy Rubio-Gonzalez, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, and Ben Liblit. EIO: Error Handling is Occasionally Correct. In *Proceedings of the 6th USENIX Symposium on File and Storage Technologies (FAST '08)*, pages 207–222, San Jose, California, February 2008.
- [20] C.A.R. Hoare. Monitors: An Operating System Structuring Construct. *Communications of the ACM*, 17(10), October 1974.
- [21] Steve Jobs, Bertrand Serlet, and Scott Forstall. Keynote Address. Apple World-wide Developers Conference, 2006.
- [22] Horatiu Jula, Daniel Tralamazza, Cristian Zamfir, and George Candea. Deadlock Immunity: Enabling Systems to Defend Against Deadlocks. In *Proceedings of the 8th Symposium on Operating Systems Design and Implementation (OSDI '08)*, San Diego, California, December 2008.
- [23] Hyojun Kim, Nitin Agrawal, and Cristian Ungureanu. Revisiting Storage for Smartphones. In *Proceedings of the 10th USENIX Symposium on File and Storage Technologies (FAST '12)*, San Jose, California, February 2012.
- [24] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Michael Norrish, Rafal Kolanski, Thomas Sewell, Harvey Tuch, and Simon Winwood. seL4: Formal Verification of an OS Kernel. In *Proceedings of the 22nd ACM Symposium on Operating Systems Principles (SOSP '09)*, Big Sky, Montana, October 2009.
- [25] Zhenmin Li, Shan Lu, Suvda Myagmar, and Yuanyuan Zhou. CP-Miner: A Tool for Finding Copy-paste and Related Bugs in Operating System Code. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation (OSDI '04)*, San Francisco, California, December 2004.
- [26] Zhenmin Li, Lin Tan, Xuanhui Wang, Shan Lu, Yuanyuan Zhou, and Chengxiang Zhai. Have Things Changed Now? – An Empirical Study of Bug Characteristics in Modern Open Source Software. In *Workshop on Architectural and System Support for Improving Software Dependability (ASID '06)*, San Jose, California, October 2006.
- [27] Zhenmin Li and Yuanyuan Zhou. PR-Miner: Automatically Extracting Implicit Programming Rules and Detecting Violations in Large Software Code. In *Proceedings of the 13th ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE '05)*, Lisbon, Portugal, September 2005.
- [28] Shan Lu, Soyeon Park, Eunsoo Seo, and Yuanyuan Zhou. Learning from Mistakes — A Comprehensive Study on Real World Concurrency Bug Characteristics. In *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XIII)*, Seattle, Washington, March 2008.
- [29] Cathy Marshall. "It's like a fire. You just have to move on": Rethinking Personal Digital Archiving. Keynote at FAST 2008, February 2008.
- [30] Chris Mason. The Btrfs Filesystem. oss.oracle.com/projects/btrfs/dist/documentation/btrfs-ukuug.pdf, September 2007.
- [31] Avantika Mathur, Mingming Cao, Suparna Bhattacharya, Alex Tomas Andreas Dilge and, and Laurent Vivier. The New Ext4 filesystem: Current Status and Future Plans. In *Ottawa Linux Symposium (OLS '07)*, Ottawa, Canada, July 2007.
- [32] Marshall K. McKusick, William N. Joy, Sam J. Leffler, and Robert S. Fabry. A Fast File System for UNIX. *ACM Transactions on Computer Systems*, 2(3):181–197, August 1984.
- [33] Marshall Kirk McKusick, William N. Joy, Samuel J. Leffler, and Robert S. Fabry. fsck - The UNIX File System

- Check Program. Unix System Manager's Manual - 4.3 BSD Virtual VAX-11 Version, April 1986.
- [34] Sean Morrissey. *iOS Forensic Analysis: for iPhone, iPad, and iPod Touch*. Apress, 2010.
 - [35] Yoann Padioleau, Julia Lawall, René Rydhof Hansen, and Gilles Muller. Documenting and Automating Collateral Evolutions in Linux Device Drivers. In *Proceedings of the EuroSys Conference (EuroSys '08)*, Glasgow, Scotland UK, March 2008.
 - [36] Nicolas Palix, Gael Thomas, Suman Saha, Christophe Calves, Julia Lawall, and Gilles Muller. Faults in Linux: Ten Years Later. In *Proceedings of the 15th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XV)*, Newport Beach, California, March 2011.
 - [37] Vijayan Prabhakaran, Lakshmi N. Bairavasundaram, Nitin Agrawal, Haryadi S. Gunawi, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. IRON File Systems. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP '05)*, pages 206–220, Brighton, United Kingdom, October 2005.
 - [38] Eric S. Raymond. *The Cathedral & the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*. O'Reilly, October 1999.
 - [39] Mendel Rosenblum and John Ousterhout. The Design and Implementation of a Log-Structured File System. *ACM Transactions on Computer Systems*, 10(1):26–52, February 1992.
 - [40] Cindy Rubio-Gonzalez, Haryadi S. Gunawi, Ben Liblit, Remzi H. Arpaci-Dusseau, and Andrea C. Arpaci-Dusseau. Error Propagation Analysis for File Systems. In *Proceedings of the ACM SIGPLAN 2009 Conference on Programming Language Design and Implementation (PLDI '09)*, Dublin, Ireland, June 2009.
 - [41] Suman Saha, Julia Lawall, and Gilles Muller. Finding Resource-Release Omission Faults in Linux. In *Workshop on Programming Languages and Operating Systems (PLOS '11)*, Cascais, Portugal, October 2011.
 - [42] Swarup Kumar Sahoo, John Criswell, and Vikram Adve. An Empirical Study of Reported Bugs in Server Software with Implications for Automated Bug Diagnosis. In *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering (ICSE '10)*, Cape Town, South Africa, May 2010.
 - [43] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. The Hadoop Distributed File System. In *Proceedings of the 26th IEEE Symposium on Mass Storage Systems and Technologies (MSST '10)*, Incline Village, Nevada, May 2010.
 - [44] Mark Sullivan and Ram Chillarege. Software Defects and their Impact on System Availability – A Study of Field Failures in Operating Systems. In *Proceedings of the 21st International Symposium on Fault-Tolerant Computing (FTCS-21)*, Montreal, Canada, June 1991.
 - [45] Mark Sullivan and Ram Chillarege. A Comparison of Software Defects in Database Management Systems and Operating Systems. In *Proceedings of the 22st International Symposium on Fault-Tolerant Computing (FTCS-22)*, pages 475–484, Boston, USA, July 1992.
 - [46] Adan Sweeney, Doug Doucette, Wei Hu, Curtis Anderson, Mike Nishimoto, and Geoff Peck. Scalability in the XFS File System. In *Proceedings of the USENIX Annual Technical Conference (USENIX '96)*, San Diego, California, January 1996.
 - [47] Stephen C. Tweedie. Journaling the Linux ext2fs File System. In *The Fourth Annual Linux Expo*, Durham, North Carolina, May 1998.
 - [48] Xi Wang, Haogang Chen, Zhihao Jia, Nickolai Zeldovich, and M. Frans Kaashoek. Improving Integer Security for Systems. In *Proceedings of the 10th Symposium on Operating Systems Design and Implementation (OSDI '12)*, Hollywood, California, October 2012.
 - [49] Yin Wang, Terence Kelly, Manjunath Kudlur, Stphane Lafortune, and Scott Mahlke. Gadara: Dynamic Deadlock Avoidance for Multithreaded Programs. In *Proceedings of the 8th Symposium on Operating Systems Design and Implementation (OSDI '08)*, San Diego, California, December 2008.
 - [50] Weiwei Xiong, Soyeon Park, Jiaqi Zhang, Yuanyuan Zhou, and Zhiqiang Ma. Ad Hoc Synchronization Considered Harmful. In *Proceedings of the 9th Symposium on Operating Systems Design and Implementation (OSDI '10)*, Vancouver, Canada, December 2010.
 - [51] Junfeng Yang, Can Sar, and Dawson Engler. EXPLODE: A Lightweight, General System for Finding Serious Storage System Errors. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI '06)*, Seattle, Washington, November 2006.
 - [52] Junfeng Yang, Paul Twohey, Dawson Engler, and Madanlal Musuvathi. Using Model Checking to Find Serious File System Errors. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation (OSDI '04)*, San Francisco, California, December 2004.
 - [53] Zuoning Yin, Ding Yuan, Yuanyuan Zhou, Shankar Pasupathy, and Lakshmi Bairavasundaram. How Do Fixes Become Bugs? – A Comprehensive Characteristic Study on Incorrect Fixes in Commercial and Open Source Operating Systems. In *Proceedings of the Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE '11)*, Szeged, Hungary, September 2011.