

Towards Automatically Checking Thousands of Failures with Micro-specifications



Haryadi S. Gunawi
Thanh Do
Pallavi Joshi
Joseph M. Hellerstein
Andrea C. Arpaci-Dusseau
Remzi H. Arpaci-Dusseau
Koushik Sen

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2010-98

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-98.html>

June 16, 2010

Copyright © 2010, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Towards Automatically Checking Thousands of Failures with Micro-specifications

Haryadi S. Gunawi, Thanh Do[†], Pallavi Joshi, Joseph M. Hellerstein,
Andrea C. Arpaci-Dusseau[†], Remzi H. Arpaci-Dusseau[†], and Koushik Sen

University of California, Berkeley

[†] University of Wisconsin, Madison

Abstract

Recent data-loss incidents have shown that existing large distributed systems are still vulnerable to failures. To improve the situation, we propose two new testing approaches: failure testing service (FTS) and declarative testing specification (DTS). FTS enables us to systematically push a system into thousands of failure scenarios, leading us to many critical recovery bugs. With DTS, we introduce “micro-specifications”, clear and concise specifications written in Datalog style, which enables developers to easily write, refine, and manage potentially hundreds of specifications.

1 Introduction

The power of large clusters behind cloud computing has brought us not only benefits but also a new challenge: a growing number and frequency of failures that must be managed [9, 10]. Failing to deal with failures will directly impact the reliability and availability of data and jobs. Unfortunately, recent data-loss incidents experienced by a telecommunication provider [16], a popular social networking site [18], and a large bank [20] still display the vulnerability of existing systems to hardware failures, failures such as machine crashes, disk and network failures. This leaves us with an important question: How can we verify the correctness of large distributed systems in dealing with a growing number failures?

We believe a proper answer requires two more rigorous approaches than the current state-of-the-art of testing. First, we must begin with a new approach that can systematically push a system into many possible failure scenarios. Some existing tools are often not equipped with a fault-injection feature, and thus failures are injected manually. Some others have the feature, but only inject failures of the same type (e.g., crashes) [25]. When it comes to injecting a wide variety of failures, the state-of-the-art in industry is to do it randomly [3, 13, 24].

Second, after thousands of failures are systematically injected, we still need to verify the correctness of many properties of the system. For a large distributed system,

there are potentially hundreds of properties that must be checked (especially under failures). However, with existing approaches, a check sometimes has to be written in tens of lines of code (e.g., in C++ [25], or a scripting language [14]). The drawbacks are two-fold. First, they hinder developers from writing a large number of specifications; in practice, the number of deployed checks for a new system is typically small, and hence does not scale to the complexity of the system. Second, even if we have hundreds of specifications (e.g., in old systems where the developers have incrementally added them over the years), they could be as big as the system code itself. As a result, the specifications are also likely to be wrong and correcting them is not always straightforward [8]. Further complicating the matter, the specifications must also evolve as the system evolve [13].

To overcome the first challenge above, we present *failure testing service* (FTS), a framework that can systematically exercise many combinations of varieties of failures (and by “many” our target is “thousands of scenarios”). The key ingredient to a systematic fault-injection technique is to identify failure points and a list of possible failures that could happen at a particular failure point (crashes, disk/network faults, etc.). A failure point and an injected failure constitute a *failure ID* that FTS will explore; exercising a combination of failures is essentially exercising a combination of failure IDs.

To enable developers to write specifications quickly and incrementally, we introduce *declarative testing specification* (DTS). At the heart of DTS is a declarative relational logic language (based on Datalog) that enables a check to be written in only a few lines. We chose Datalog style due to its nature of compactness and expressiveness in specifying logical relationships [1, 7, 15, 21]. With DTS, we promote a practical iterative style of writing specification: a developer can begin with a handful of high-level specifications, and as she finds new bugs, she can add more detailed specifications to precisely pinpoint the bugs. If the same bugs appear again in the future (as the code being developed), the tighter specifications will allow the developer to avoid wasting another hours of debugging time.




Figure 1: **HDFS Write Protocol.** The figure illustrates when and where failures could occur in this protocol. Communications to the master node is excluded for brevity.

This paper makes the three following contributions:

- We introduce failure IDs as a new abstraction for exploring failures systematically. In addition, we also provide a ready-to-use “failure surface” for Java-based systems (Section 2).
- With DTS, we introduce “micro-specifications”, clear and concise specifications written in Datalog style, which enables developers to manage hundreds of specifications. We also present a design pattern for declarative debugging, that is, how a developer could refine specifications iteratively (Section 3).
- We have applied FTS and DTS to the Hadoop File System (HDFS), an open-source version of the Google File System [5], and uncovered many critical recovery bugs (Section 4). We chose HDFS as our first target due to its complexity (over 25 KLOC) and its growing popularity (it has been widely deployed in over 80 medium to large organizations including Amazon, Yahoo, and Facebook).

After presenting our contributions, we follow with a discussion of new challenges (Section 5) and close with related work (Section 6) and conclusion (Section 7).

2 Failure Testing Service

Figure 1 motivates the need for a systematic failure service: in a distributed write, there are many points where the system components could fail (labeled with X). The expectation is that the write protocol should succeed if at least one datanode is alive [5]. The developer might want to verify that the specification holds even if, for example, the disk at the second datanode fails in the first round; or, if the first datanode *and* the second datanode crash in the second round.

To help developers in this regard, the ultimate goal of FTS is to exercise as many combinations of failures as possible. In a sense, this is similar to model checking which explores different sequences of states. One key




Figure 2: **FTS Framework.**

technique in model-checking is to record the states that have been explored, which is commonly done by hashing the abstract state of the system and recording the history of the hashes. Similarly in our case, we need a clear abstract representation of a failure, which can be hashed and thus recorded in the failure history. Below, we first introduce the concept of failure IDs and follow with the overall framework. So far, we have used FTS to exercise more than 40,000 combinations of failures.

2.1 Failure IDs

To construct a failure ID, we first define a *failure point*, a system/library call that performs disk or network I/O. For every failure point, FTS generates a list of possible failures that could happen before and after. For example, FTS could throw an exception before a disk-write failure point. FTS could also crash a node after the node receives a message but before it sends an acknowledgment. When FTS exercises a failure ID (*i.e.*, injects a particular failure at a particular failure point), FTS records the hash value of the failure ID in the failure history.

Table 1 shows an example of a failure ID. The table also shows that a failure point contains more complex information (static, dynamic, and domain-specific) than what we have described above. These information are essential to increase failure coverage. For example, if a static failure point could be called by different upper-level functions, then dynamic information such as stack trace is useful to expand the failure scenarios. Domain-specific information such as target I/O is valuable because a common function could write to different file types or send messages to different nodes.

2.2 FTS Framework

We plan to integrate the concept of failure IDs to a model checker (since a failure ID does not contain the system state). However, since not all systems support a ready-to-use model-checker, we build our own testing framework that could be quickly integrated to the system we want to test. Figure 2 depicts the overall framework of FTS. We first instrument the system (*e.g.*, HDFS) by inserting a “failure surface” which builds failure IDs and send

Info Type	Field	Example value
Static	Func. call	: OutputStream.flush()
	I/O type	: Write
	Location	: BlockReceiver.java
	Line	: 45
Dynamic	Stack trace	: (the stack trace)
Domain specific	Thread name	: Block Receiver
	Target I/O	: Disk1, Metafile
FTS	Failure	: Bad disk
	Hash Value	: 1289065658

Table 1: **A Sample of Failure ID.** A failure ID contains static, dynamic, and domain-specific information about a failure point and an injected failure. Hash is used to label a failure ID.

them to the FTS server, which then makes failure decisions based on the failure history. The developer attaches the workload to be tested at the workload driver and specifies the maximum number of failures (MAX). The while loop exits when the server does not see a new combination of MAX failure IDs.

2.2.1 Failure Surface

There are many possible layers to introduce a failure surface (e.g., inside the library or at the VMM layer). We do this between the application and the Java library for two reasons. First, at this layer, we could get rich domain-specific information useful to increase failure coverage. Second, by leveraging aspect-oriented programming (AspectJ in our case), this can be done without any modification to the system. Since our target audience is system developers, we assume source code availability.

The failure surface has two important jobs. First, it builds the non-domain and domain specific information of every failure ID (Table 1). The non-domain specific part is directly available from AspectJ. For the domain-specific part, sometimes not enough information is available at the failure point. For example, given a call to `DataOutputStream.write()`, the target output stream (e.g., what the disk file or the network address is) is not available. Thus, we have written and applied a context-passing aspect to the Java SDK so that relevant contexts are accessible at the failure point. This is a general feature that can be used for all Java-based systems.

Second, when the FTS server returns a failure confirmation, the failure surface must map the failure to the correct error-code (if the failure is non-crash). For Java-based systems, this is straightforward: a failure point either returns false or throws an I/O exception. For C-based systems, a failure point could return a variety of error-codes (e.g., EIO, EROFS), and hence more mappings are needed.

2.2.2 FTS Server and Workload Driver

One inherent complexity lies in the interaction between the driver and the server. In essence, we do not know what failure IDs that server will see before we run the workload. Yet, we want to somewhat “control” the combinations of failure IDs for a given run. As a more concrete example, consider three failure IDs A, B, and C, not known apriori. For two failures, we want to inject {A, B} in one run, {A, C} in another run, and so on. To control this, we implement *failure locking*: after the first run, the driver will lock the first failure to A (`lock[1] = A`) such that in the next run the server could only inject A for the first failure. For the second failure, since the driver does not lock it (`lock[2] = ∅`), the server will inject *any* new failure that it has not exercised (e.g., C). Eventually, the driver will unlock the first failure in order for the server to exercise {B, C}.

2.2.3 Other Features

We briefly discuss other important features of FTS below. First, failure ID should not contain non-deterministic information (e.g., chunk filename [5]), otherwise FTS always gets a new hash failure ID and will never finish. Thus, we have a domain-specific component that filters out irrelevant non-deterministic information. Second, since some failures (e.g., bad disks) are only emulated, the server keeps track persistent failures such that future related operations will naturally fail (similar to persistent network partitioning in [24]). Finally, FTS server supports extensible filters to exercise specific failures in specific contexts. This feature is useful to reproduce the sequence of failures that leads to the bug.

2.3 Summary

We have introduced failure IDs as a new abstraction for exploring failures. The ultimate goal is to exercise as many combinations of failure IDs as possible. In a sense, this is similar to model checking which exhausts all abstract states of a system. Thus, to have a model-checker exercise failures systematically, the notion of failure IDs can be directly employed.

Second, we have built a ready-to-use FTS. Deploying FTS is relatively easy; one could quickly do that without the domain-specific component. For example, we have ported FTS to two other systems in just a few hours (Section 4.4). To increase failure coverage, one can gradually introduce the domain-specific component.

FTS is written in 4180 lines of code in Java. As of now, FTS is able to inject transient failures, persistent disk failures, crashes, and disk corruption. In the future, we will support other types of failures (e.g., delays to emulate message reordering).

3 Declarative Testing Specification

After failures are injected, the developer needs to verify the system correctness. This can be done in many ways (e.g., via external behaviors or detailed internal checks). For example, one can easily observe the external behavior (e.g., return value of the client API, hangs, or crashes). A harder but more precise way is to write detailed checks. With the state of the art, a single check could reach one hundred lines of code [8, 13, 24]. As a result, managing hundreds of checks could be complicated (not to mention that they must evolve as the system evolves [13]).

With DTS, we attempt to improve the state-of-the-art of writing testing specification. To achieve this, we explored writing “micro-specifications”, clear and concise specifications written in Datalog style. We chose this style as it has been successfully used in building distributed systems declaratively [1, 15] and in verifying some aspects of system correctness (e.g., security [7, 21]). However, since the innovation of high-performance declarative language is still underway [1], we feel that using Datalog for just writing system specifications is a sweet spot to explore.

In the next two sections, we show how micro-specifications promote a practical iterative style of writing specifications. More specifically, the first example shows how a developer can *refine* loose specifications into tighter ones, while the second example illustrates how to *incrementally add* more detailed specifications. The motivation for these examples is that a developer never begins with complete and precise specifications. But, as she unearths new problems, she might wish to refine existing specifications or add more specifications to pinpoint similar problems quickly in the future.

3.1 Refining Specifications

This section focuses on demonstrating how a developer could refine a high-level specification into a tighter one. Here, we use the HDFS log recovery process as an example. We begin with a high-level specification that will catch data-loss bugs:

```
lostFiles(F) :- userFiles(F), not-in server(F)
```

The specification is written as a Datalog rule which consists of a head (`lostFiles`) and predicates in the body (`userFiles` and `server`). The head is evaluated when the body is true. A comma between predicates represents conjunction. Thus, the rule specifies the relation: “a user file `F` is lost if it does not exist at the server.”

Next, we need to “fill” the test predicates with runtime facts; the head and the predicates are essentially database tables. The `userFiles` table is filled whenever the HDFS client write API returns a success. Populating the `server` table is a bit more complicated. As a background, HDFS

maintains two server files that store the metadata of user files: the image file (`img`) and the log file (`log`). When a client stores a file, it is first recorded in the log file. Upon a master reboot (or periodically), the log recovery process will merge the two files into the image file and empty the log file. To ensure idempotency, HDFS utilizes two other files (`log2` and `img2`). In short, *at any time*, user files should exist in the union of all the four files. Thus, we can simply write the specification below, which reads: “a user file `F` exists at the server if `F` exists in any of the four server files.” Uppercase and lowercase letters (e.g., `F` and `img`) stand for variables and constants respectively.

```
server(F) :- filesIn(F, img);
server(F) :- filesIn(F, log);
server(F) :- filesIn(F, log2);
server(F) :- filesIn(F, img2);
```

We use FTS to automatically insert all possible crashes in this process. Interestingly, the rules above trigger a violation of the data-loss specification. However, since the rules are not rigorous enough, we were not able to exactly pinpoint the bug. We spent a couple hours debugging, and then refined the specification (and hence the iterative process of writing specifications).

The new specification reflects in detail the five-step process shown in Table 2. That is, depending on the progress, user files are expected to be in a different subset of the four files as shown in the last column of the table. For example, during the process of merging `img` and `log` to `img2` at step 3, user files are expected to be in `img` and `log`, but not in `img2`, because `img2` is still not complete. However, when `log` is emptied at step 4, `img2` is complete, and thus user files should be found in `img2`, but not in `img` and `log`. The way HDFS keeps track the steps is via the existence of the four files (column 1). To express this, we simply introduce two new relational tables `exists(img,log,log2,img2)` and `step(Num)`:

```
step(1) :- exists(1, 1, 0, 0);
step(2) :- exists(1, 1, 1, 0);
step(3) :- exists(1, 1, 1, 1);
step(4) :- exists(1, 1, 0, 1);
```

Using the new relations above, we can simply write new `server` rules that correspond to the logic in the last column of Table 2, that is, “user files should be in `img2` at step 4, or in `img` and `log` otherwise”:

```
server(F) :- filesIn(F, img2), step(N), N == 4;
server(F) :- filesIn(F, img), step(N), N != 4;
server(F) :- filesIn(F, log), step(N), N != 4;
```

With this more rigorous specification, we were able to pinpoint the data-loss bug. As a background, if the master node crashes exactly before the log renaming operation at step 4, the protocol will begin again from step 1 (not shown). However, exactly before step 4, there is a bug

img	log	log2	img2	Steps:	User files (f_1 and f_2) in:
f_1	f_2	—	—	1. Start	img + log
f_1	f_2	\emptyset	—	2. Create an empty log2	img + log
f_1	f_2	\emptyset	$f_1 f_2$	3. Merge img and log to img2	img + log
f_1	\emptyset	—	$f_1 f_2$	4. Rename log2 to log	img2
$f_1 f_2$	\emptyset	—	—	5. Rename img2 to img	img + log

Table 2: **Log Recovery Protocol.** The table shows the process of merging user-file metadata (f_1 and f_2) in `img` and `log`. If a crash occurs in the middle, the next reboot will start from a different step depending on which of the four files exist (not shown). \emptyset and — represent empty and non-existent file respectively.

that truncates the log file (f_2 is no longer in `log`). Thus, if a crash happens after the truncate bug and before step 4, the next reboot will merge `img` (f_1) with an empty `log`! The final `img` file will only contain f_1 , and hence f_2 is lost. This bug still exists in the latest released version of HDFS.

The `lostFiles` rule captures the bug because all rules that depend on `log` are re-evaluated when the bug truncates `log`: file f_2 is removed from `filesIn(F, log)`, which will then remove f_2 from `server(F)` (because the bug is at step 3). Finally, f_2 is added to `lostFiles` because f_2 is in `userFiles` but not in `server`. In summary, in this section, we have written a full-page prose specification of the log recovery process, and it has been elegantly summarized in just 8 lines of Datalog rules.

3.2 Adding Specifications Incrementally

To illustrate the generality of the iterative process above, we briefly give another example where we catch unavailability bugs. This time, rather than refining a specification, we incrementally add more specifications to pinpoint the bugs. We begin with a general specification below which reads: “there is an unavailability bug if a write operation fails but there is at least one good datanode.”

```
unavailability(F, Num) :-
  write_op(F, fail), good_nodes(Num), Num >= 1;
```

We caught some bugs that break the rule. One bug is a buggy failover where the master keeps returning the same datanodes that the client cannot reach. This is because it takes the master 10 minutes to detect a dead datanode by default. But, since the master does not incorporate the client’s view of unreachable datanodes, the master does not give the client another set of available datanodes. To catch this specific design bug, we simply added a new rule:

```
bad_failover(BadNode) :-
  master_returns(BadNode), unreachable_nodes(BadNode);
```

3.3 Implementation

To fill the rules with runtime facts, one can write some form of state-exposer that interposes the internal functions

of each node [13]. However, we decided to interpose and reverse-engineer the system API and the disk and network protocols, primarily because inter-node protocols and on-disk formats change less frequently than the internal functions. The obtained information is then “translated” into Datalog. For example, when a buffer is written to the `img` file, we extract the file entries in the buffer, and add them to `filesIn(F, img)`; when the HDFS write API returns a success, the corresponding file is added to `userFiles(F)`; when HDFS creates and deletes any of the four storage files, the `exists()` relation is re-evaluated. After the translation phase, the tester enters the declarative world in which she can build the abstract model of the system succinctly. The DTS translation mechanism for HDFS is written in 1200 lines of code in Java.

4 Preliminary Results

We have integrated FTS and DTS to HDFS ver. 0.20.0. We chose HDFS as our first target due to its complexity (over 25 KLOC) and its growing popularity (it has been widely deployed in over 80 medium to large organizations including Amazon, Yahoo, and Facebook). To show the effectiveness of FTS and DTS, we present our preliminary results in four aspects: Can FTS exercise thousands of failure scenarios? Does FTS help us find critical bugs? Can we write hundreds of specifications with DTS? Finally, is our framework useful for other systems as well?

4.1 Thousands of Failures

We have run FTS on three HDFS workloads: write, append and master reboot. For every workload, we use different setups: we vary the resources (datanodes and disks) and the number of injected failures. The high-level expectation is that as long as there is at least one resource available, the workload should not fail. Given a workload, FTS will run a number of experiments until there is no new set of failures to inject. As an example, for a write workload with 3 replicas, 3 available datanodes, and 2 bad-disk failures per run, FTS identifies 74 failure IDs and exercises 2415 combinations of two failures. Because of the expo-

ponential growth of the possible combinations (Section 5), three is the maximum number of failures we could inject per run.

In total, FTS has generated 394 unique failure IDs and exercised 41332 unique combinations (with a maximum of three failures per run), out of which 22872 result in unsuccessful experiments. The ability to explore this large number of combinations has led us to numerous critical recovery bugs.

We also found that injecting multiple failures is crucial: we found some sections of code that would not have been executed unless three failures are injected. More specifically, we have exercised 36 failure IDs that only appear as a third failure, 8 out of which result in unsuccessful experiments. With FTS, the vision of “towards 100% coverage of recovery code” becomes feasible.

4.2 Bugs Found and Observations

We are still in the process of debugging the thousands of unsuccessful experiments. As of now, we have uncovered 21 bugs (15 are new) that lead to unavailability, data loss, and inconsistencies. We present our high-level observations below.

First, unsurprisingly, an implementation is much more complex than the prose specification; when we read the simple prose specification of the HDFS write protocol, we had little confidence that we would find new bugs. But, we were wrong; we found many bad failover strategies that fail the corresponding operations even though good resources are still available. Second, replicas are only useful to prevent a single point of hardware failure, but software is still the single point of failure. A prime example is the data-loss bug described in Section 3.1; the bug causes the master node to lose all replicas. Another example is some bugs in the append protocol that make the datanodes lose all replicas of the affected files.

4.3 Hundreds of Specifications

So far, we have been writing specifications in a bug-driven manner, that is, for every bug that we found (from external observation), we wrote more specifications to pinpoint the bug more precisely. Overall, we have written 50 rules in only 70 lines. After writing these rules, we uncovered new bugs not observable externally (hence the power of detailed checks). We believe that there is nothing that prevents us from writing many more specifications, which is our next top priority.

4.4 Other Systems

To show that FTS can be deployed to other systems quickly, we have ported FTS to two different versions

of the Apache Lucene concurrency library in just a few hours. In this particular experiment, we would like to find new concurrency bugs in the presence of failures (*e.g.*, a hang bug because an exception block does not wake up waiting threads); we note that many novel concurrency bug-finding tools often do not incorporate failures. We have run 23 workloads (out of the available 184 workloads in their JUnit tests) and found 4 concurrency bugs (2 are new).

5 Future Work and Discussion

We have begun deploying FTS and DTS to two other large distributed systems: Cassandra (an open-source version of Amazon Dynamo and Google BigTable) and ZooKeeper (an open-source version of Google Chubby).

Apart from this initial success, there are some challenges that we need to solve. The first relates to non-determinism; we found a bug that only appears if the parallel node executions run in a certain order. Thus, we need some form of execution scheduling [19]. We also plan to add other failure types such as message reordering and network partitioning.

The second challenge is to intelligently prune the failure points; the number of combinations of failures grows exponentially (*e.g.*, in a particular setup with 149 failure IDs and 3 injected failures per experiment, FTS generates 17263 thousands of experiments in over five hours). This prohibits us to explore more failures per experiment. We have sketched out some pruning techniques, but unfortunately they are out of the scope of this paper.

Finally, we need some form of heuristics to classify unsuccessful experiments to the actual bugs [6]. From our experience, thousands of unsuccessful experiments map to only a few number of bugs. Without heuristics, we have to debug one experiment at a time.

6 Related Work

Throughout the paper, we have compared our work with others. In this section, we highlight the distinction again.

One state-of-the-art of finding bugs is to use a model-checker. Most model-checkers rely on the abstract state of the system [11, 23, 25, 24]. Although, some of them have built-in “knobs” to enable/disable failures, it is still unclear how they exercise combinations of failures. To do that systematically, these model checkers could use our concept of failure ID as part of the abstracted state.

LFI [17] presents an XML-based language to specify the to-be-injected failures. Their language can be useful for FTS (*e.g.*, for writing a filter to reproduce an unsuccessful experiment).

Existing model checkers are equipped with checkers. WiDS and D3S use their own scripting language to write checks [14, 13]. In this language, a check often reaches more than 50 lines. We found that writing checks in relational logic language is more clear and concise. FiSC and EXPLODE use existing file system checkers written in C [25, 26]. Again, reasoning about hundreds of non-declarative checks is proven to be hard [8]. D3 uses a variant of Datalog by writing checks on top of distributed log messages [4], while ours run on top of the interposed disk and network communications. Singth *et al.* also uses a variant of Datalog, but their checks only run on systems built on the same language [22]. Cloud9 also plans to design a new testing specification that is accessible to programmers [2].

Another way to build robust systems is via writing an executable specification in theorem prover, as demonstrated by seL4 [12]. In a sense, their work and our work share a similarity: we write two versions of the system in different languages; they use Isabelle/HOP, while we use Datalog style.

7 Conclusion

As failure becomes the norm, when data is lost or availability is reduced, we should no longer blame the failure, but rather the inability to recover from the failure. Thus, we began with a simple question: How can we assess the quality of existing systems in handling thousands of possible failures? This simple question gave birth to FTS. As we achieve this end, we were confronted with the fact that, for a large distributed system, there are a large number of properties that need to be verified (not just state invariants, but also specific behaviors of the recovery protocols). This gave birth to DTS, which has enabled us to write, refine, and manage many micro-specifications easily; the more specifications we can write, the more properties we can verify, the finer bug reports we will produce, and the less time we will spend on debugging.

References

- [1] Peter Alvaro, Tyson Condie, Neil Conway, Khaled Elmeleegy, Joseph M. Hellerstein, and Russell C. Sears. BOOM Analytics: Exploring Data-Centric, Declarative Programming for the Cloud. In *Proceedings of the EuroSys Conference (EuroSys '10)*, Paris, France, April 2010.
- [2] George Candea, Stefan Bucur, and Cristian Zamfir. Automated Software Testing as a Service. In *Proceedings of the 2010 ACM Symposium on Cloud Computing (SoCC '10)*, Indianapolis, Indiana, June 2010.
- [3] Tushar Chandra, Robert Griesemer, and Joshua Redstone. Paxos Made Live - An Engineering Perspective. In *Proceedings of the 26th ACM Symposium on Principles of Distributed Computing (PODC '07)*, Portland, Oregon, August 2007.
- [4] Byung-Gon Chun, Kuang Chen, Gunho Lee, Randy H. Katz, and Scott Shenker. D3: Declarative Distributed Debugging. UC Berkeley Technical Report No. UCB/EECS-2008-27, 2008.
- [5] 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.
- [6] Kirk Glerum, Kinshuman Kinshumann, Steve Greenberg, Gabriel Aul, Vince Orgovan, Greg Nichols, David Grant, Gretchen Loihle, and Galen Hunt. Debugging in the (Very) Large: Ten Years of Implementation and Experience. In *Proceedings of the 22nd ACM Symposium on Operating Systems Principles (SOSP '09)*, Big Sky, Montana, October 2009.
- [7] Salvatore Guarneri and Benjamin Livshits. Gatekeeper: Mostly Static Enforcement of Security and Reliability Policies for JavaScript Code. In *Proceedings of the 13th USENIX Security Symposium (Sec '09)*, Montreal, Canada, August 2009.
- [8] Haryadi S. Gunawi, Abhishek Rajimwale, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. SQCK: A Declarative File System Checker. In *Proceedings of the 8th Symposium on Operating Systems Design and Implementation (OSDI '08)*, San Diego, California, December 2008.
- [9] James Hamilton. On Designing and Deploying Internet-Scale Services. In *Proceedings of the 21st Large Installation System Administration Conference (LISA '07)*, Dallas, Texas, November 2007.
- [10] Alyssa Henry. Cloud Storage FUD: Failure and Uncertainty and Durability. In *Proceedings of the 7th USENIX Symposium on File and Storage Technologies (FAST '09)*, San Francisco, California, February 2009.
- [11] Charles Killian, James Anderson, Ranjit Jhala, and Amin Vahdat. Life, Death, and the Critical Transition: Finding Liveness Bugs in Systems Code. In *Proceedings of the 4th Symposium on Networked Systems Design and Implementation (NSDI '07)*, Cambridge, Massachusetts, April 2007.
- [12] Gerwin Klein, June Andronick, Kevin Elphinstone, Gernot Heiser, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, 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.
- [13] Xuezheng Liu, Zhenyu Guo, Xi Wang, Feibo Chen, Xiaochen Lian, Jian Tang, Ming Wu, M. Frans Kaashoek, and Zheng Zhang. D3S: Debugging Deployed Distributed Systems. In *Proceedings of the 5th Symposium on Networked Systems Design and Implementation (NSDI '08)*, San Francisco, California, April 2008.

- [14] Xuezheng Liu, Wei Lin, Aimin Pan, and Zheng Zhang. WiDS Checker: Combating Bugs in Distributed Systems. In *Proceedings of the 4th Symposium on Networked Systems Design and Implementation (NSDI '07)*, Cambridge, Massachusetts, April 2007.
- [15] Boon Thau Loo, Tyson Condie, Joseph M. Hellerstein, Petros Maniatis, Timothy Roscoe, and Ion Stoica. Implementing Declarative Overlays. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP '05)*, Brighton, United Kingdom, October 2005.
- [16] Om Malik. When the Cloud Fails: T-Mobile, Microsoft Lose Sidekick Customer Data. <http://gigaom.com>.
- [17] Paul Marinescu and George Candea. LFI: A Practical and General Library-Level Fault Injector. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN '09)*, Lisbon, Portugal, June 2009.
- [18] Lucas Mearian. Facebook temporarily loses more than 10% of photos in hard drive failure. www.computerworld.com.
- [19] Madanlal Musuvathi, Shaz Qadeer, Tom Ball, Gerard Basler, Piramanayakam Arumuga Nainar, and Iulian Neamtiu. Finding and Reproducing Heisenbugs in Concurrent Programs. In *Proceedings of the 8th Symposium on Operating Systems Design and Implementation (OSDI '08)*, San Diego, California, December 2008.
- [20] John Oates. Bank fined 3 millions pound sterling for data loss, still not taking it seriously. www.theregister.co.uk/2009/07/22/fsa_hsbcc_data_loss.
- [21] Xinming Ou, Sudhakar Govindavajhala, and Andrew W. Appel. MulVAL: A logic-based network security analyzer. In *Proceedings of the 14th USENIX Security Symposium (Sec '05)*, Baltimore, Maryland, July 2005.
- [22] Atul Singh, Petros Maniatis, Timothy Roscoe, and Peter Druschel. Using Queries for Distributed Monitoring and Forensics. In *Proceedings of the EuroSys Conference (EuroSys '06)*, Leuven, Belgium, April 2006.
- [23] Maysam Yabandeh, Nikola Knezevic, Dejan Kostic, and Viktor Kuncak. CrystalBall: Predicting and Preventing Inconsistencies in Deployed, Distributed Systems. In *Proceedings of the 6th Symposium on Networked Systems Design and Implementation (NSDI '09)*, Boston, Massachusetts, April 2009.
- [24] Junfeng Yang, Tisheng Chen, Ming Wu, Zhilei Xu, Xuezheng Liu, Haoxiang Lin, Mao Yang, Fan Long, Lintao Zhang, and Lidong Zhou. MODIST: Transparent Model Checking of Unmodified Distributed Systems. In *Proceedings of the 6th Symposium on Networked Systems Design and Implementation (NSDI '09)*, Boston, Massachusetts, April 2009.
- [25] 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.
- [26] 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.