

Composable Security Analysis of OS Services*

Ran Canetti^{2,**}, Suresh Chari¹, Shai Halevi¹, Birgit Pfizmann¹,
Arnab Roy¹, Michael Steiner¹, and Wietse Venema¹

¹ IBM T.J. Watson Research Center

² Tel-Aviv University

Abstract. We provide an analytical framework for basic integrity properties of file systems, namely the binding of files to filenames and writing capabilities. A salient feature of our modeling and analysis is that it is *composable*: In spite of the fact that we analyze the filesystem in isolation, security is guaranteed even when the file system operates as a component within an arbitrary, and potentially adversarial system.

Our results are obtained by adapting the *Universally Composable* (UC) security framework to the analysis of software systems. Originally developed for cryptographic protocols, the UC framework allows the analysis of simple components in isolation, and provides assurance that these components maintain their behavior when combined in a large system, potentially under adversarial conditions.

1 Introduction

Contemporary software systems are complex, consisting of many millions of lines of code, spread across a myriad of components and sub-components. A natural approach for analyzing such large systems is by analyzing each component separately, and “hoping” to use the component-wise analysis to analyze the entire system. Unfortunately, applying this approach to security analysis is problematic. Even if a component is simple enough to analyze separately, its interaction with other components can yield unexpected results. Often, a component will be used in environments different from what its designers initially had in mind, alongside other components that perhaps did not even exist when the original component was analyzed, potentially violating some assumptions that were made in the analysis.

Ideally, we would like to analyze the behavior of a component in isolation, and have the assurance that this behavior remains intact *even when that component is embedded in a new environment*. Within the realm of cryptography, the frameworks of Reactive Simulatability [19,1] and Universal Composability (UC) [4,5] ensure just that. These frameworks are aimed at capturing the security of cryptographic primitives and protocols, ranging from authentication and key exchange,

* This work is supported in part by the Department of Homeland Security under the grant FA8750-08-2-0091.

** Supported by the Check Point Institute for Information Security. Part of this work was done when visiting IBM T.J. Watson Research Center.

to public-key encryption and signatures, zero-knowledge, and more (see [5] for many examples.) However, many of the features of these frameworks appear at first to be specific to the realm of cryptographic protocols. A natural question is whether the “composable security” approach sketched above can be carried out in a meaningful way even outside the limited domain of cryptography. A positive answer could significantly reduce the overhead in analyzing the security of large systems, while at the same time provide better overall security guarantees.

In this work we demonstrate that this can indeed be done, in the context of guaranteeing some basic integrity properties of filesystems. For this purpose we adapt the UC framework to software systems by establishing new conventions for modeling process management and scheduling. The current work is one of just a few attempts to apply the UC formalism to a large and complex software system, and we believe that it will enable further application of the UC formalism to other software systems.

Analysis in the UC framework proceeds by defining an idealized specification model and an implementation, and then proving that the implementation realizes the idealized specification. Our main contribution is a very simple filesystem specification model, called SIMPFS, that captures many integrity concerns in contemporary filesystems, together with an implementation over existing POSIX filesystems [20] and a proof that the implementation realizes the specification model.

The composability properties of our analysis imply that software systems that use our implementation over POSIX behave essentially the same as if they were using the simple, idealized specification system SIMPFS.

This is a very strong security guarantee. In particular, it allows analyzing software systems without worrying about how the filesystem is implemented, and without worrying about potential bad interactions between the analyzed system and the filesystem implementation.

Our filesystem model is geared toward ensuring integrity of files and their names, and in particular preventing filename manipulation attacks. In such attacks, a victim program expects a particular filename to have certain semantics. (E.g., a mail program may expect the file `/var/mail/root` to be the mail file for the super-user.) In the attack, the adversary creates a link by the same name in the filesystem, pointing to another file (e.g., `/var/mail/root -> /etc/passwd`), thereby “tricking” the victim program into accessing an unexpected file. (In the mail example, such a link may cause a naive mail program to write incoming email into the system’s password file.) Such attacks were quite common in UNIX systems of old.

Our implementation builds on the ideas presented in Chari et al. [7], who address the problem of privilege escalation attacks via filename manipulation. To counter these attacks, Chari et al. present a “safe” name resolution procedure, and deploy this system-wide on popular POSIX systems. The SIMPFS interfaces are designed to tightly bind files with their names: files can be accessed *only* via the names they were created with, which means that filename manipulation attacks are impossible in our model. Our proof — showing that implementation

based on [7] realizes the model — implies in particular that it indeed eliminates these filename manipulation attacks.

SIMPFS offers a simple interface that captures enough filesystem primitives for application developers to build meaningful applications. The simplicity of SIMPFS is due to its very narrow interface (only four commands) and the fact that *it does not have directories*. We argue that the murky relation between files and their names in plain POSIX systems stem to a large extent from the fact that pathnames consist of many directories, each with its own permissions, which are combined in a non-obvious manner to yield the effective permissions for the entire name. In contrast, a filename in SIMPFS is just a single entity with explicitly specified permissions. Thus SIMPFS provide applications with radically simplified semantics, making it easier to use the filesystem without falling into traps. At the same time, we argue that the vast majority of contemporary applications in POSIX systems *do not really need directories*, and can be implemented over the simple SIMPFS interface without loss of functionality.

1.1 Related Work

Triggered by Joshi and Holzmann’s mini-challenge [13], there is a lot of recent work on formalization and verifications of file systems. Most notably, Freitas et al [10] specify and prove a POSIX file store in Z/Eves. This body of work focuses mostly on the correctness aspects and does not address in depth the security and access control aspects of filesystems. In the broader perspective of (secure) operating systems, there is a long history of formalization and verification, from PSOS [16] to the recent seL4 [14]. While they make considerable progress toward high-assurance OS, these works are not based on frameworks that allow easy composition of components to form larger systems. Additionally, the focus in many of these works is on mandatory access control whereas we cover a discretionary control. (We stress that although our model addresses integrity concerns, these are very different from the Biba integrity model [3].)

An abstract model of another large standard systems, the browser, suitable for proofs of cryptographic protocols exists in [12]. It includes a model of information-flow properties under attack. However, the federated identity protocols built on top of it have only been proven secure with respect to specific security properties, not in a real-world / ideal-world setting [12].

Protocol Composition Logic (PCL) [8] is a comparable general approach on reasoning about (cryptographic) network protocols in a composable fashion. Recently, PCL was applied to analyze systems [9], more specifically integrity properties provided by TPM. The symbolic and axiomatic nature of PCL leads to a more axiomatic specification of security rather than the declarative form in UC. Furthermore, the composition theorems in PCL are weaker than in the UC framework.

A noteworthy contribution to secure composition of large systems is the CHATS project [17], that identifies architectural principles to guide the structuring and decomposition of trustworthy systems. That work is largely orthogonal to ours, as it does not focus on formal modeling or proofs.

There have been many more attempts to leverage well-established formalisms such as logic, typing or process calculi to model composability of certain system security properties, e.g., McLean [15] for non-interference properties or Bengtson et al [2] for cryptographic protocols and access control mechanisms. Many of them provide tool support; but they do not provide the same composition guarantees as in the UC framework.

Organization. Due to lack of space, we did not include in this writeup the proof of security in the UC framework. A full version that includes the proof is available online [6]. In Section 2 we outline the UC framework, then in Section 3 we describe our SIMPFS model and in Section 4 we show how it is implemented over POSIX.

2 The Universal Composability Framework

We briefly describe the relevant aspects of the framework of universally composable (UC) security. The reader is referred to [4] for more details. The framework describes two probabilistic games: The *real world* that captures the protocol flows and the capabilities of an attacker, and the *ideal world* that captures what we think of as a secure system. The notion of security asserts that these two worlds are essentially equivalent.

THE REAL-WORLD MODEL. The players in the real-world model are all the entities of interest in the system (e.g., the nodes in a network, the processes in a software system, etc.), as well as *the adversary* A and *the environment* \mathcal{Z} . All these players are modeled as efficient, probabilistic, message-driven programs (formally, they are all interactive Turing machines).

The actions in this game should capture all the interfaces that the various participants can utilize in an actual deployment of this component in the real world. In particular, the capabilities of A should capture all the interfaces that a real-life attacker can utilize in an attack on the system. (For example, A can typically see and modify network traffic.) The environment \mathcal{Z} is responsible for providing all the inputs to the players and getting all the outputs back from them. Also, \mathcal{Z} is in general allowed to communicate with the adversary A . (This captures potential interactions where higher-level protocols may leak things to the adversary, etc.)

THE IDEAL-WORLD MODEL. Security in the UC framework is specified via an “ideal functionality” (usually denoted \mathcal{F}), which is thought of as a piece of code to be run by a completely trusted entity in the ideal world. The specification of \mathcal{F} codifies the security properties of the component at hand. Formally, the ideal-world model has the same environment as the real-world model, but we pretend that there is a completely trusted party (called “the functionality”), which is performing all the tasks that are required of the protocol. In the ideal world, participants just give their inputs to the functionality \mathcal{F} , which produces the correct outputs (based on the specification) and hands them back to the participants. \mathcal{F} may interact with an adversary, but only to the extent that

the intended security allows. (E.g., it can “leak” to the adversary things that should be publicly available, such as public keys.) Specifying the code of \mathcal{F} is typically a non-trivial task. It is important that \mathcal{F} satisfies all the desired security properties, but also that \mathcal{F} does not impose unnecessary constraints: It is only too easy to write a functionality that describes “what we intuitively want”, but is not realizable by any implementation. Another crucial concern is the *simplicity* of the functionality \mathcal{F} , since we want \mathcal{F} to capture the important security concerns, not the mundane implementation details.

UC-SECURITY AND THE COMPOSITION THEOREM. An implementation π **securely realizes** an ideal functionality \mathcal{F} if no external environment can distinguish between running the protocol π in the real world and interacting with the trusted entity running the ideal functionality \mathcal{F} in the ideal world. That is, for every adversary A in the real world, there should exist an adversary A' in the ideal world, such that no environment \mathcal{Z} can distinguish between interacting with A and π in the real world and interacting with A' and \mathcal{F} in the ideal world.

The striking feature of the UC framework is its ability to handle composition. Specifically, the composition theorem from [4] asserts the following: Let ρ be an arbitrary system that runs in the ideal world and uses (perhaps multiple copies of) the functionality \mathcal{F} . Next, consider the system ρ' in the real world, that is the same as ρ except that in ρ' each call to the ideal functionality \mathcal{F} is replaced by executing the implementation π . Then, if π securely realizes \mathcal{F} it is guaranteed that system ρ' behaves essentially the same as system ρ . In particular, all the security properties of ρ are inherited by protocol ρ' . This guarantee is the basis for the composable security guarantees provided by the UC framework.

2.1 Conventions for Software Systems

We briefly describe some technicalities that must be resolved when attempting to apply the UC framework to software system, and the conventions that we use to address them. The “entities of interest” in our work are processes, which differ somewhat from the interactive Turing machines (ITMs) in common cryptographic models. One aspect relates to side-channels: whereas an ITM can only influence other ITMs by sending messages, a process shares some physical resources with other processes on the same machine, so it could influence them via side channels such as timing and concurrency. In this work we ignore that aspect, i.e. we do not have any side channels in our formal model. (This does not matter for our current SIMPFS model, since we do not model any secrecy requirements.) We thus just let the adversary learn “whatever it needs,” so it has no use for side channels.

A more important difference is preemptive multitasking: common crypto models postulate a sequential scheduling model, where an active ITM keeps the control until it sends a message, at which point the recipient becomes active. On the other hand, processes in contemporary OSes can be made to yield control involuntarily. Resolving this discrepancy is not as hard as it may seem, since (side-channels aside) an active entity has no effect on its surroundings until it sends a message, which means that influencing the surroundings only comes with

losing the control. We use the standard sequential scheduling of the UC framework, but ensure that the adversary gets the control after every message is sent, and can decide when this message will be delivered. (This is somewhat similar to the “buffer scheduler” from [1].) Hence the adversary in our formal model is able to simulate the actions that would have happened in the actual deployed system, delay delivery messages until the simulation arrives at the point where they were delivered. We thus argue that the formal adversary in our model is able to induce any behavior that can happen in the actual deployed system.

Another difference is that some processing in real systems is done not by the processes themselves, but by the kernel on their behalf. Hence also in our model we postulate the existence of a “kernel component” that can do things on behalf of processes. In our filesystem example, this kernel component is only responsible for maintaining the process privileges: Whenever a process calls a filesystem function, the kernel adds the process-id and roles of the calling process to the list of arguments, and forwards everything to the filesystem. (The kernel component gets these roles from the environment.) We note that although we do not use it in our filesystem example, in general we could have several such “kernel components” in a system, representing several physical machines.

3 SimpFS: A Simple Idealized File-System

This section describes SIMPFS, our simple filesystem model. SIMPFS has a minimalistic interface with simple semantics, having only basic primitives to create, read, write and delete files. Still, we believe that the this file-system functionality is sufficient for most applications. (Other aspects — such as locking — can be implemented on top of our interface.) The SIMPFS model includes file write permissions, hence capturing properties of *filesystem integrity*. We currently do not model read permissions, but we expect that this work can be extended to include read permissions without too much change.

An important feature of SIMPFS is that *it does not have any directories, only files and their names*. As we mention in the introduction, we believe that directories have “inherently cumbersome semantics”, hence decided to do away with them in order to keep the semantics as simple as possible. We stress that the model supports names that include ‘/’ (so applications can still store their temporary data in files with names that begin with “/tmp/”). But a name such as “/a/b/foo” is viewed as just one entity, and its existence does not imply the existence of an object with name “/a/b.” Of course, our *implementation* over POSIX still interprets ‘/’ as a directory separator, and name creation induces the right associations between names and paths, in spite of symlinks, adversarial write permissions etc. While directories are a useful and convenient way to manage and organize systems, we argue that directory permissions are very rarely needed in applications (if ever), and most applications can therefore directly use the SIMPFS interface.

A key security property of SIMPFS is that it rules out filename manipulation attacks. Our focus on this property is motivated by the large number of privilege escalation attacks due to unsafe pathname resolution that were discovered in

POSIX systems over the years. A classical example of this type of attacks is local mail delivery, where `/var/mail` may be world-writable, allowing an adversary to create a link from `/var/mail/root` to (say) `/etc/passwd`, thereby “tricking” a naive mail-delivery program (running as `root`) to write the content of incoming mail into `/etc/passwd`. Such attacks arise due to the opaque mapping of names to files in POSIX. SIMPFS features a very tight binding between files and their names: a file can be manipulated *only* with the names it was created with.

We describe an implementation of SIMPFS over contemporary POSIX filesystems and *rigorously prove* that this implementation realizes SIMPFS, using the UC framework. The proof implies that processes that use our implementation will be protected against pathname manipulation attacks such as above even if adversarial processes use the same POSIX filesystem in arbitrary ways.

3.1 A Formal Model of SimpFS

SIMPFS consists of files and their names. A newly created file is given some names, and thereafter the file can be accessed by any of these names. Existing names can be deleted, but one cannot add names to existing files. When deleting names, a file can end up with zero names, in which case it is not reachable anymore so we can consider it as deleted. We associate permissions with both the file names and the files themselves:

- Every file has a list of roles that can write in it, called the *Writers* list. A process can write to a file if it holds a role in the Writers list of the file.
- File names have a set of *Manipulators*, listing all the roles that have permission to delete that name.

In the current version we do not have read permissions, which means that SIMPFS allows every process to read every file.

In more details, our ideal SIMPFS maintains an array of files and an associative array of names: `files []` is an array of files (indexed by integers). Each entry is a file, consisting of an array of bytes (i.e., a data blob) and a list of roles (specifying the Writers of this file). `names []` is an associative array (indexed by strings). We refer to the index of an entry as a file-name, and each entry consists of a pointer to a file (i.e., an integer) and a list of roles (specifying the Manipulators of this name). The interface below constrains the Manipulator lists, making sure that all the names of the same file have the same set of Manipulators. (This choice is not very important, it is done mostly to simplify the presentation.)

In the initial state, the file-system is empty, with no files and no names (i.e., both arrays are empty). There are only four operations that are supported in SIMPFS: `CreateFile` creates a new file with some names, `DeleteName` deletes an existing name, `Read` reads data from a file (specified by some name), and `Write` writes data to a file (specified by some name).

The semantics of these operations is described by the pseudo-code in Figure 1. As is the case with every formal UC functionality, the pseudo-code includes not only the intended functionality as seen by the legitimate users of the system, but also all the interfaces that an adversary can utilize to attack it. This is codified

```

CreateFile(Writers, Manipulators, Names, pid, Roles)
{
  // Allow the adversary to fail the operation and decide the error code
  var retCode = AdversaryAction("CreateFile",Writers,Manipulators,Names,pid,Roles);
  if (retCode != OKAY) return retCode;

  var codes[] = empty;    // a local list of return codes, one per name
  var f = index of next available entry in the files[] array;
  files[f].data=empty, files[f].Writers=Writers;

  // Allow the adversary to decide whether to create each name
  for each fName in Names {
    var code = AdversaryAction("CreateOneName", fName);
    if (code!=OKAY) codes[i]=code;
    else {
      if (names[fName] already exists) codes[i] = FILE_EXISTS;
      else {
        names[fName].file=f, names[fName].Manipulators=Manipulators;
        codes[i]=OKAY;
      } } }
  call AdversaryAction("Done CreateFile") and then return codes;
}

DeleteName(fName, pid, Roles)
{
  // Allow the adversary to fail the operation and decide the error code
  var retCode = AdversaryAction("DeleteName",fName,pid,Roles);
  if (retCode != OKAY) return retCode;

  if (names[fName] does not exist) return FILE_DOESNT_EXIST;
  if (Roles intersect names[fName].Manipulators = emptyset) return NO_PERMISSION;

  delete names[fName]; // Note: no point deleting the file, even if not reachable
  call AdversaryAction("Done DeleteName") and then return OKAY;
}

Write(fName, atAddr, data, pid, Roles)
{
  // Allow the adversary to fail the operation
  var retCode = AdversaryAction("OpenWrite",fName,pid,Roles);
  if (retCode != OKAY) return retCode;

  if (names[fName] does not exist) return FILE_DOESNT_EXIST;
  var f = names[fName].file;    // f serves as a "handle" to the file
  if (Roles intersect files[f].Writers = emptyset) return NO_PERMISSION;

  var numBytes = AdversaryAction("Write",fName,atAddr,data,pid,Roles);
  if (numBytes < length(data)) truncate data to numBytes bytes;    // only partial write

  var nBytes = length(data);
  if (atAddr < 0) atAddr = length(files[f].data);    // append
  else if (atAddr > length(files[f].data)) {
    prepend (atAddr-length(files[f].data)) zero bytes to data;
    atAddr = length(files[f].data);
  }
  write data to files[f].data starting at position atAddr;
  call AdversaryAction("Done Write") and then return [OKAY,nBytes];
}

Read(fName, fromAddr, nBytes, pid, Roles)
{
  // Allow the adversary to fail the operation or read less bytes
  var [retCode,numBytes] = AdversaryAction("Read",fName,fromAddr,nBytes,pid,Roles);
  if (retCode != OKAY) return retCode;
  if (numBytes < nBytes) nBytes = numBytes;

  if (names[fName] does not exist) return FILE_DOESNT_EXIST;
  var f = names[fName].file;

  if (fromAddr < 0) fromAddr = 0;
  else if (fromAddr > length(files[f].data)) {
    fromAddr = length(files[f].data);
    nBytes = 0;
  }
  if (nBytes < 0)    // read to end-of-file
    nBytes = length(files[f].data) - fromAddr;

  data = content of files[f].data from fromAddr for nBytes;

  call AdversaryAction("Done Read") and then return [OKAY,nBytes,data];
}

```

Fig. 1. The SIMPFS commands

by an `AdversaryAction` call, in which SIMPFS “leaks” to the adversary the details of its operation, and also lets the adversary influence these operations.

A key feature of SIMPFS is that a file can be accessed *only* using one of the names that were specified when the file was created, thus eliminating filename-manipulation attacks such as described above. Hence proving that an implementation realizes SIMPFS implies in particular that such attacks cannot be successfully mounted against the implementation.

We make no liveness guarantees in SIMPFS, so at the beginning of every operation the adversary is given the option to abort the operation and determine the error code. (This does not mean that an implementation of SIMPFS cannot ensure some liveness properties, but it means that a proof that an implementation realizes SIMPFS carries no such guarantees within itself.)

The pseudo-code includes with every call also the process-id and permissions (Roles) of the caller, which in our system model are filled by the kernel component. (Formally there is also an implicit “invocation id” for each call of one of the four main operations, allowing SIMPFS to handle messages received from the ideal-world adversary for different invocations.) Note also that the `AdversaryAction` at the beginning and end of every operation comply with our convention that the adversary gets the control before any message is delivered. Finally, we note that all the variables in the code in Figure 1 are local to that invocation, except for the global `files[]` and `names[]`.

PROCESS CORRUPTION. Following the standard conventions of the UC framework, SIMPFS has a special procedure to handle the case where the adversary corrupts a process. For our purposes it is more convenient to let the environment decide when a process is corrupted (as opposed to the adversary, which is the more common convention in UC-model works). When the environment corrupts a process, this process makes a call `IamCorrupted(pid, Roles)`, to inform SIMPFS that “it belongs to the adversary” now. SIMPFS informs the adversary of this call, and it remembers that this process and all its roles are now bad. Thereafter, the adversary is allowed to make all the usual calls to SIMPFS (`CreateFile`, `DeleteName`, `Read`, `Write`) on behalf of that process. SIMPFS will process these calls just as if it was the corrupted process that made the call, but will return the result to the adversary rather than to the environment.

Every call from the corrupted process (not via the adversary) will be routed directly to the adversary, and the adversary can always instruct SIMPFS to send anything to the corrupted process (which will then be forwarded to the environment). Also, if the roles of the corrupted players change then the kernel component will notify SIMPFS of this change. SIMPFS will add any new role that a corrupted process acquires to its list of bad roles, but *it will not remove any roles from that list*, even if the corrupted process loses some of its roles. (This last aspect represents the fact that the corrupted process may already have used this role to introduce artifacts into the filesystem, that will remain even after the process no longer has this role.)

ATOMICITY OF THE SIMPFS OPERATIONS. The operations `DeleteName` and `Read` are atomic, whereas `CreateFile` and `Write` are not: In `DeleteName` and

Read, once the adversary allows the operation to go through (by returning `OKAY`), `SIMPFS` holds onto the control-flow throughout the name lookup and the operation itself, and only then it yields control back to the adversary.

In **Write**, on the other hand, the control is returned to the adversary after the file lookup (via the call `AdversaryAction("Write", ...)`), and only then is the operation carried out. Similarly in **CreateFile**, the adversary gets the control before the creation of any name. This choice was made so that we would be able to realize `SIMPFS` over the POSIX interface that requires to open the file and then write in it. The real-world read can be made atomic by checking after the fact that the file did not change since it was opened, but for write such a check is meaningless since the file was already written.

MAPPING UNIX PERMISSIONS TO ROLES. The interfaces of `SIMPFS` above are defined with “generic roles” that encode permissions, with access control being a simple role inclusion. Our implementation over POSIX, of course, uses userids and groups, which are particular types of roles. The mapping is quite straightforward, roughly there is a different role for each userid and group in the system, and a process gets the role corresponding to its effective-uid and all the roles corresponding to its groups. There is also one role for “others”, that every process has. Some care must be taken since POSIX permissions do not exactly follow role inclusion. (For example, if a file is not owner-readable then the owner cannot read it, even if the file is readable by “others”.) Adjusting the mapping to this technicality is quite straightforward, and is omitted here.

4 Implementing `SimpFS` over POSIX

We describe `simpfs`, which is a concrete implementation of the `SIMPFS` functionality over the POSIX filesystem interface [20]. The presentation below focuses on a user-space implementation, where each `simpfs` operation runs with the effective uid of its caller, but we point out that the same procedures can also be implemented in the kernel.

Our implementation relies on the “safe pathname resolution” procedure of Chari et al. [7], that protects processes from opening adversarial links. While resolving paths this procedure ensures that an adversary can not manipulate the resolution to result in opening unintended components. In `simpfs`, very roughly speaking, each operation consists of first using that procedure to open the corresponding file and then performing the actual operation.

Before describing this implementation, we first introduce concepts that are used in the rest of the paper and describe some assumptions that we make on the POSIX filesystems underlying our implementation. Then in Section 4.2 we describe the `safeDirOpen` procedure, which is the heart of our implementation and builds on [7], and then in Section 4.3 we describe the rest of the implementation.

4.1 Concepts and Properties of POSIX

We assume that the reader is familiar with basic concepts of POSIX such as directories, pathnames, users and groups, hardlinks and symlinks, etc.

Definition 1 (Pathname Manipulators). *Let $/dir1/.../dirn/foo$ be an absolute pathname. The manipulators of this pathname are all the roles (users and groups) that own, or have write permissions in, any directory visited during the resolution of this pathname.*

Note that the definition applies even when a pathname does not resolve, and that `root` is a manipulator of every pathname.

Definition 2 (Safe Names). *A pathname is system safe if its only manipulator is `root`. A pathname is safe for U (where U is a user-id) if its only manipulators are `root` and U . Otherwise, the pathname is unsafe for U .*

For example, in a typical UNIX system the pathname `/etc/passwd` is system safe, the pathname `/home/joe/mbox` is safe for user `joe`, and the pathname `/var/spool/mail/jane` is unsafe for everyone (as `/var/spool/mail` may be world- or group-writable).

Definition 3 (Simple Pathnames). *A pathname is simple if it is an absolute path that resolves to a regular file, its elements are only hard links (i.e., not symbolic links), no elements are named `.` or `..`, and the pathname contains no repeated slashes `//`.*

ASSUMPTIONS. We now list some properties that we assume on the underlying POSIX system, and use in our proof of security. Most of these assumptions are justified either by the fact that they are part of the POSIX specification itself, or by the fact that many contemporary POSIX filesystems seem to satisfy them.

Assumption 1. *The underlying filesystem does not contain multiple mount points to the same filesystem, and each directory has only one parent (i.e., one hard link with a name other than `.` or `..`).*

Justification. Assumption 1 is justified by the fact that nearly all contemporary POSIX implementations either do not allow processes to create additional hard links to directories (e.g., FreeBSD, Linux) or restrict this operation to the super-user (e.g., Solaris, HP-UX). A notable exception is MacOS.

We observe that given Assumption 1, for every reachable hard link to a regular file there is a unique simple name that ends with that hard link. Moreover a resolution of any absolute name that ends with that hard link will visit all the directories in this unique simple pathname.

Assumption 2 (Permissions) 1. *If an operation by a process affects the content of a file, then the process must have write permission for that file. 2.* *Let P be an absolute pathname. If an operation by a process affects the resolution of P or changes the permissions or ownership of any of the directories visited during its resolution, then that process must have a role which is a manipulator of P .*

Justification. The only operations that affect pathname resolution are creating, removing, or renaming pathname components, and they all require write permission in the containing directory. Also, note that only the owner of a directory (or `root`) can change the permissions of that directory, and in most systems only `root` can change ownership.

Corollary 1. *Let P be some pathname, denote by $\mathcal{M}(P)$ the set of manipulators for P (user-ids and groups), and let \mathcal{B} be a set of roles such that $\mathcal{M}(P) \cap \mathcal{B} \neq \emptyset$. Then changing the manipulator set for P so that $\mathcal{M}(P) \cap \mathcal{B} = \emptyset$ requires an operation by a process with some role outside of \mathcal{B} .*

This corollary follows from the fact that no role can perform an operation to “remove itself” from the manipulators set of P (which follows from Assumption 2). Details are deferred to the full version.

Assumption 3. *The hardlink to a directory in its parent directory can only be removed when the child directory is empty. Moreover, after the hardlink is removed from the parent directory, no further entries can be created in the child directory, even if some process still holds a handle to it.*

Justification. The last part of Assumption 3 is justified by the fact that `rmdir` implementations remove the entries ‘.’ and ‘..’ from the child directory before removing the hard link in the parent directory, and no new entries can be created in directories without ‘.’ and ‘..’.

Corollary 2. *If a system call for creating an entry in a directory returns successfully, then the hard link for this directory in its parent directory could not have been removed before that system call, or removed after the call but before the newly-created entry is removed.*

4.2 The `safeDirOpen` Procedure

Underlying our `simpfs` implementation is a procedure for *safe name resolution*, which is adapted from the work of Chari et al. [7]. Our `safeDirOpen` procedure takes an absolute pathname, resolves it “in a safe manner” and returns a handle to the directory containing the final hard link to the actual file, the name of that hard link, and additional information as discussed below. The top-level operations of `simpfs` first call `safeDirOpen` and then perform the requested operation on the final hard link.

`safeDirOpen` resolves a pathname one atom at a time, each time opening the next atom (or reading it, if it is a symlink), while keeping track of the owners and writers of the visited directories. (Below we identify the time that a directory was visited as the time when it was opened, and the time that a symlink was visited with the time that it was read.)

The procedure can be in one of three states: system-safe, safe-for-uid, or unsafe. When invoked (by a process with effective uid U), the procedure begins in a system-safe state, switching to safe-for-uid state upon visiting a directory where U is an owner or writer, and switching to unsafe state upon visiting a directory with any writer or owner other than `root` or U . Once in unsafe state it stays in that state for the duration of the current name resolution. Likewise, there is no transition from safe-for-uid to the system-safe state.

When `safeDirOpen` enters the unsafe state, it does not follow symlinks for the remainder of the current name resolution. Also, for technical reasons the

procedure never accepts pathnames that contain multiple slashes ‘//’ or have components named ‘.’ or ‘..’, and it refuses to visit any directory whose name begins with the special prefix `_SimpFS_ephemeral_`. In any of these cases, the procedure returns an error code.

Once `safeDirOpen` arrives at the final atom (and verifies that it is indeed the final atom and not a symlink), it ends successfully, returning a handle to the directory containing this last hard link, as well as the name of the hard link. In addition, `safeDirOpen` returns its current state (system-safe, safe-for-uid, or unsafe), the set of owners and writers of the directories that it visited, and an array of (handle,name) pairs, containing handles to all visited directories, and the names that were looked-up in those directories. (These names could belong to either a directory, a symlink, or the final hard link.)

Upon failure, `safeDirOpen` returns an error code, a handle to the last directory pathname component that was successfully resolved, the state (system-safe, etc.) and manipulators of that directory, and the unresolved remainder of the pathname. For example, when called to resolve `/a/b/c`, if it encountered an error after visiting `/a` but before visiting `/a/b`, then it will return a handle to directory `/a`, the state and manipulators of `/a`, and the remainder of the pathname argument “`b/c`”. (Note that this will be the return value even if `/a/b` happens to be a symlink and the procedure visited more directories after `/a`, but could not completely resolve `/a/b`.)

4.3 Implementing the `simpfs` Commands

`createFile(Writers,Manipulators,Names)`. When called by a process with effective-uid U , the procedure begins by checking that U belongs to the set of manipulators specified by the `Manipulators` parameter. Then it creates a new file with an ephemeral name that begins with the special prefix `_SimpFS_ephemeral_`. This ephemeral name is created so that it is safe for U , thus ensuring that no other users can remove or rename it.¹

Now `createFile` attempts to set the write permissions of the new file as specified in the `Writers` parameter. If this is successful, it proceeds to create the names, one at a time, by calling the subroutine `createOneName` for each name in `Names`. After all the calls to `createOneName`, the procedure `createFile` removes the ephemeral name that it created for the new file, and returns the vector of return codes that it received from all the calls to `createOneName`.

The subroutine `createOneName(fName)` begins by checking that the new name is an absolute name, and that it does not contain ‘//’ or elements named ‘.’ or ‘..’, or elements that begin with `_SimpFS_ephemeral_`. Then it calls `safeDirOpen(fName)` thus obtaining a handle to the last successfully resolved directory on this pathname and the corresponding set of manipulators. If all the directories were resolved successfully, then `createOneName` checks that the set of manipulators equals the `Manipulators` parameter, and aborts if they differ.

If some directories were not resolved, `createOneName` verifies that the manipulator set of the prefix is not too large (i.e., it must be contained in the

¹ See Section 4.5 for a short discussion of this point.

`Manipulators` parameter), aborting otherwise. Then `createOneName` attempts to create the remaining directories, one at the time, initially creating each one so that it is only writable by owner U with an ephemeral name that begins with `_SimpFS_ephemeral_`. Upon success, it tries to set the write permissions of the last directory so that the resulting set of manipulators will match the `Manipulators` parameter. Then it goes over all the newly created directories, top to bottom, renaming each one to the name that it is supposed to have according to `fName`.

Once all the directories exist and have the right set of manipulators and the right names, the procedure `createOneName` makes a `linkat` system call to create a hard link in the last directory, pointing to the new file. `createOneName` then returns whatever code was returned from the `linkat` system call.

If any operation fails, then `createOneName` attempts to clean-up after itself, trying to remove all the directories that still have names that begin with `_SimpFS_ephemeral_`. However, after a directory was renamed to its “permanent name”, `createOneName` does not remove it.

In the proof of security in the full version we rely on the following properties of our implementation of `createFile`:

- The initial ephemeral name for the new file is safe for the effective-uid of the calling process.
- The procedure never creates symlinks, only directories and hard links.
- The procedure only changes permissions and/or removes pathname components if these components begin with the special prefix `_SimpFS_ephemeral_`.
- A name `fName` is created if and only if the `linkat` system call at the end of the subroutine `createOneName(fName)` is successful.

`deleteName(fName)`. When called with effective-uid U , `deleteName` calls `safeDirOpen(fName)` and aborts if that function fails. Else `deleteName` has an array of pairs (handle,name), and the state with which `safeDirOpen` arrived at the final directory (system-safe, safe-for-uid, or unsafe). If the state is not system-safe, then `deleteName` checks that the final directory is either world-writable, or owner-writable and owned by U , and it aborts otherwise.² Also, if the state is unsafe then `deleteName` checks that the file that the hard link points to has only a single hard link, aborting otherwise.

Then `deleteName` attempts to delete the final hard link, followed by attempts to delete the directories higher-up on the path. `deleteName` returns when any system call to remove a name fails, or when any of these names resolves to a symlink, or when it is done deleting all the names in the array. The return code from `deleteName` is whatever was returned from the first `unlink` system call (i.e., the one that deleted the hard link at the end of `fName`).

We note that barring a race condition, this implementation of `deleteName` does not delete symlinks. In the proof in the full version we show that the only

² This check is intended to protect against privilege-escalation attacks on `setgid` programs, cf. Section 4.5.

cases where these race conditions are possible are when the adversary already has permissions to delete these symlinks by itself.

`read(fName,...)`. When called with effective-uid U , `read` calls `safeDirOpen(fName)` to get a handle for the final directory, the name of the hard link pointing to the actual file, and the state at which it arrived in this last directory: `system-safe`, `safe-for-uid`, or `unsafe`. Then `read` uses `openat`, `lstatat` and `fstat` to open the file and verify that it is still the same file (and not a symlink). In addition, if the state is not `system-safe`, then `read` checks that the file is either world-readable, or owner-readable and owned by U , and it aborts otherwise. Also, if the state is `unsafe` then `read` checks that the file has only a single hard link, aborting otherwise.

Then the procedure uses the `read` system call to read the file, and before closing the file it makes yet another `lstatat` system call to check that the hard link still points to the same inode as it did when it was `opened`. If all these checks pass, then `read` returns the result from the `read` system call.

`write(fName,...)`. The procedure `write` is almost identical to `read` except that it adds a write-permission check on the actual file, and it does not do the final check after writing to verify that the hard link still points to the same inode. (Indeed, such check is useless since the file was already written to.)

4.4 Consistency Properties of the Implementation

In the proof of security in the full version, it is important to consider what changes may happen in the filesystem between the time that the `safeDirOpen` pathname resolver visits some directory and the time that the procedure that called `safeDirOpen` returns. An important technical observation is that if the procedure that called `safeDirOpen` was successful then none of those visited directories could have been removed during this time.

Lemma 4. *Consider an execution of one of the procedures `createOneName`, `deleteName`, `read`, or `write` on argument $fName$, and assume that the procedure succeeds (i.e., does not return an error code). Assume further that no symlink that was read during name resolution was later deleted or renamed during the execution of this procedure, and no directory was renamed after it was `opened` by this procedure. Then also none of these directories was deleted after it was `opened` and before the time that the procedure issued the system call (respectively, `linkat`, `unlinkat` or `openat`) for the final hard link in $fName$.*

Moreover, for the procedures `createOneName`, `read`, and `write`, as long as no symlinks are deleted or renamed, no directories are renamed, and the final hard link in $fName$ exists in its original containing directory, then also none of these directories is deleted even after the operation returns.

The proof follows from Assumption 3, and is provided in the full version. \square
 Jumping ahead, we use Lemma 4 in the proof by noting that our SIMPFS implementation never renames or removes symlinks, or renames directories, and hence no uncorrupted process will do any of these things. If in addition we know that no corrupted process has write permissions in any of the directories visited

then also corrupted processes could not rename or remove symlinks or rename directories. Thus, we can apply Lemma 4 and conclude that all the directories stay put throughout the execution of `createOneName`, `deleteName`, `read`, or `write`.

4.5 Rationale and Discussion

Before proceeding to the formal proof of security, we discuss here some of the rationale for our implementation, including some specific attacks that the implementation was designed to foil.

Privilege-escalation attacks on `setgid` programs. Our implementation of `safeDirOpen` only considers the effective-uid for the purpose of determining the safety of a directory, and thus we must consider the possibility of privilege-escalation attacks between processes with the same effective-uid. In contemporary UNIX systems, two processes with the same effective-uid can have different filesystem privileges only if one of them has a group-privilege that the other does not,³ as would happen when one of these processes runs a `setgid` program.

To see the problem, consider two processes running with effective-uid of `joe`, one having the additional group privilege of `mail` while the other is compromised by an attacker (e.g., due to a buffer-overflow vulnerability). Ideally, we would like to argue that files which have read/write permissions for the `mail` group (but not user `joe`) are still protected against the compromised process.

Assume that the non-compromised process with `mail` group privileges needs to delete a file `/home/joe/dir/foo`. The compromised process can create a symlink `/home/joe/dir -> /var/mail`, “tricking” the other process into deleting `/var/mail/foo` (assuming that `/var/mail/` is writable by group `mail`). Embedding this attack in our formal model, we have a name `/var/mail/foo` for which `joe` is *not a manipulator*, and a good process that attempts to delete an unrelated name `/home/joe/dir/foo`, and yet by some action of a compromised process with `joe` privileges, this results in the deletion of `/var/mail/foo`.

We fix this problem by adding a check to the operations `deleteName`, `read`, and `write`, aborting if the name is not system-safe and group privileges are needed to perform the operation. Very roughly, this defense works because it prevents the use of group privileges after following symlinks that were created by non-`root` processes. (We note that we do not need this extra precaution in `createOneName`. This is because the SIMPFS functionality restricts deletion of existing names, but puts no restrictions on the creation of names that do not exist.)

Another reason for the check of the final hard link after a `read` system call in `read` is a potential attack on open-then-read programs, which is described in the full version.

Our treatment of symbolic links. Our proof of security in relies in places on the assumption that good processes do not create symlinks. This is consistent with our `simpfs` implementation (that indeed does not create symlinks), but it begs the question why we allow `safeDirOpen` to follow symlinks at all.

³ We ignore the `fsuid` of Linux here.

The reason is that the implementation of `simpfs` is useful also in situations where the filesystem includes non-adversarial symlinks. A close inspection of our proof shows that the arguments remain valid also in the presence of non-adversarial symlinks, as long as the files that have non-adversarial symlinks in their names remain static (i.e., they are not deleted, removed, or moved). It is even possible to modify the semantics of `SIMPFS` to accommodate non-adversarial symlinks in a dynamic filesystem, but the new semantics will not be as simple anymore.

Using the sticky bit. Recall that the initial ephemeral name for a new file must be safe for the effective-uid of the calling process (denoted U). Such a name can perhaps be created in U 's home directory, but not all uid's have one. A simple way of achieving the same result in contemporary UNIX systems is creating this ephemeral name in `/tmp`, relying on the fact that `/tmp` is owned by `root` and has the sticky bit on. This does not quite fit into our definition of "safe for U " (since `/tmp` is world-writable), but it suffices for the purpose of our proof of security. Specifically, what we need is to ensure that as long as the calling process holds a handle to the new file, only U or `root` can change the resolution of the ephemeral name.

5 Conclusion

In this work we adapted the Universal Composability (UC) framework to the modeling of large software systems. Focusing on filesystem interfaces, we described `SIMPFS`, which is a simple filesystem abstraction intended to capture *filesystem integrity* concerns. We describe an implementation of this abstraction over real POSIX filesystems and prove that the implementation realizes the `SIMPFS` abstraction in the UC sense. `SIMPFS` is a simple but useful interface and with a few small enhancements is sufficient to build real applications.

Our work demonstrates that formal security frameworks such as Universal Composability can be used also beyond the niche of cryptographic protocols. Our modeling of POSIX-based file systems is the first example of this scale.

References

1. Backes, M., Pfitzmann, B., Waidner, M.: The reactive simulatability (RSIM) framework for asynchronous systems. *Inf. Comput.* 205(12), 1685–1720 (2007)
2. Bengtson, J., Bhargavan, K., Fournet, C., Gordon, A., Maffei, S.: Refinement types for secure implementations. In: 21st IEEE Computer Security Foundations Symposium (CSFS), pp. 17–32 (2008)
3. Biba, K.: Integrity considerations for secure computer systems, MITRE TR-3153, Bedford, MA (1977)
4. Canetti, R.: Universally Composable Security: A New Paradigm for Cryptographic Protocols. In: *FOCS*, pp. 136–145 (2001)
5. Canetti, R.: Security and Composition of Cryptographic Protocols. *SIGACT News* 37(3&4) (2006)

6. Canetti, R., Chari, S., Halevi, S., Pfitzmann, B., Roy, A., Steiner, M., Venema, W.: Composable Security Analysis of OS Services, Cryptology ePrint Archive, Report 2010/213, <http://eprint.iacr.org/>
7. Chari, S., Halevi, S., Venema, W.: Where Do You Want to Go Today? Escalating Privileges by Pathname Manipulation. In: Proc. Symposium on Network and Distributed Systems Security, NDSS (2010)
8. Datta, A., Derek, A., Mitchell, J.C., Roy, A.: Protocol composition logic (PCL). *Electronic Notes in Theoretical Computer Science* (2007)
9. Datta, A., Franklin, J., Garg, D., Kaynar, D.: A logic of secure systems and its application to trusted computing. In: Proc. of the IEEE Symp. on Research in Security & Privacy, pp. 221–236 (2009)
10. Freitas, L., Woodcock, J., Fu, Z.: POSIX file store in Z/Eves. *Science of Computer Programming* 74(4), 238–257 (2009)
11. Goldreich, O.: *Foundations of Cryptography*, vol. 1, 2. Cambridge Press, Cambridge (2004)
12. Gross, T.R., Pfitzmann, B., Sadeghi, A.-R.: Browser model for security analysis of browser-based protocols. In: di Vimercati, S.d.C., Syverson, P.F., Gollmann, D. (eds.) *ESORICS 2005*. LNCS, vol. 3679, pp. 489–508. Springer, Heidelberg (2005)
13. Joshi, R., Holzmann, G.J.: A mini challenge: Build a verifiable filesystem. *Formal Aspects of Computing* 19(2), 269–272 (2007)
14. Klein, G., Elphinstone, K., Heiser, G., Andronick, J., Cock, D., Derrin, P., Elkaduwe, D., Engelhardt, K., Kolanski, R., Norrish, M., Sewell, T., Tuch, H., Winwood, S.: seL4: Formal verification of an OS kernel. In: *SOSP 2009*, pp. 207–220 (2009)
15. McLean, J.: A general theory of composition for a class of ”possibilistic” properties. *IEEE Transactions on Software Engineering* 22(1), 53–67 (1996)
16. Neumann, P., Feiertag, R.: PSOS revisited. In: Omondi, A.R., Sedukhin, S.G. (eds.) *ACSAC 2003*. LNCS, vol. 2823, pp. 208–216. Springer, Heidelberg (2003)
17. Neumann, P.: Principled assuredly trustworthy composable architectures, DARPA Project CHATS, Final Rep. (2004), <http://www.csl.sri.com/users/neumann/chats.html>
18. Pfitzmann, B., Waidner, M.: A General Framework for Formal Notions of “Secure” Systems. Institut für Informatik, Hildesheim University (April 1994)
19. Pfitzmann, B., Waidner, M.: Composition and Integrity Preservation of Secure Reactive Systems. In: Proc. CCS, pp. 245–254 (2000)
20. The Open Group Base Spec. Issue 7, IEEE Std 1003.1-2008, <http://www.opengroup.org/onlinepubs/9699919799/>