Model Checking An Entire Linux Distribution for Security Violations

Benjamin Schwarz Hao Chen David Wagner
{bschwarz, hchen, daw}@cs.berkeley.edu
Geoff Morrison Jacob West
{gmorrison, jwest}@fortifysoftware.com
Jeremy Lin (jjlin@ocf.berkeley.edu) Wei Tu (tuwei@berkeley.edu)
University of California, Berkeley

Abstract

Software model checking has become a popular tool for verifying programs' behavior. Recent results suggest that it is viable for finding and eradicating security bugs quickly. However, even state-of-the-art model checkers are limited in use when they report an overwhelming number of false positives, or when their lengthy running time dwarfs other software development processes. In this paper we report our experiences with software model checking for security properties on an extremely large scale—an entire Linux distribution consisting of 839 packages and 60 million lines of code. To date, we have discovered 108 exploitable bugs. Our results indicate that model checking can be both a feasible and integral part of the software development process.

1 Introduction

Software bugs are frequent sources of security vulnerabilities. Moreover, they can be incredibly difficult to track down. Automated detection of possible security violations has become a quickly-expanding area, due in part to the advent of model checking tools that can analyze millions of lines of code [5]. One may argue that the bottleneck is no longer the computer, which runs the model checker, but rather the human, who specifies secure and insecure behaviors of programs and identifies bugs based on the output from the model checker. For this reason, model checkers that are intended to be used on real-world programs on a large scale should allow users to describe properties easily and should report errors concisely.

In this paper we describe our experience using MOPS, a static analyzer, to verify security properties in an entire Linux distribution. We use the following recipe for finding security bugs: identify an important class of security vulnerabilities, specify a temporal safety property expressing the condition when programs are free of this class of bugs, and use MOPS to decide which programs violate the property. We have developed six security properties—expressed as finite state automata (FSAs)—and refined them to minimize false positives while preserving high effectiveness. These properties aim at finding security bugs that arise from the misuse of system calls, often vulnerable interaction among these calls. For example, time-of-check-to-time-of-use (TOCTTOU) bugs involve a sequence of two or more system calls acting on the same file (see Section 3.1).

Our primary contribution is the scale of our experiment. We ran MOPS on the entire Red Hat Linux 9 distribution, which contains 839 packages totaling 60.0 million lines of code (counting total lines in all `.h`, `.c`, and `.cc` files). MOPS successfully analyzed 87% (732) of these packages; the remaining 107 packages could not be analyzed because MOPS's parser cannot parse some files in these packages[1]. To the best of our knowledge, our experiment is the largest security audit of software using automated tools reported in the literature. Model checking at this scale introduces major challenges in error reporting, build integration, and scalability. Many of these technical challenges have been addressed in our work; we show how to surmount them, and demonstrate that model checking is feasible and effective even for very large software systems.

As part of this experiment, we have worked out how to express several new security properties in a form that can be readily model checked by existing tools. Earlier work developed simple versions of some of these properties [5], but in the process of applying them at scale we discovered that major revisions and refinements were necessary to capture the full breadth of programming idioms seen in the wild. Some of the properties checked in this paper are novel; for instance, we introduce a TOCTTOU property that turned out to be very effective in finding bugs. In our experiments,

[1] The parser failed on 73 packages because they contain C++ code (the parser currently parses only C programs, although we are developing a new C++ parser), and on 34 packages because they contain certain C99 constructs.

we focused on finding bugs rather than proving their absence. Verification is difficult, especially since MOPS is not completely sound because it does not yet analyze function pointers and signals. However, we expect that our techniques could point the way to formal verification of the absence of certain classes of bugs, as better model checkers are developed in the future.

The technical contributions of this paper are threefold: 1) We show how to express six important security properties in a form that can be model checked by off-the-shelf tools; 2) We report on practical experience with model checking at a very large scale, and demonstrate for the first time that these approaches are feasible and useful; 3) We measure the effectiveness of MOPS on a very large corpus of code, characterizing the false positive and bug detection rates for different classes of security bugs.

Section 2 provides an overview of MOPS and the challenges in handling such a large number of diverse programs. Section 3 describes our security properties and the bugs that we have found in checking these properties. Section 4 reviews related work in model checking, software verification, and security. We conclude in Section 5. MOPS is freely available from `mopscode.sourceforge.net`.

2 The MOPS Model Checker

MOPS is a static (compile time) analysis tool that model checks whether programs violate security properties [6]. Given a security property—expressed as a finite-state automaton (FSA) by the user—and the source code for a program, MOPS determines whether any execution path through the program might violate the security property.

In more detail, the MOPS process works as follows. First, the user identifies a set of security-relevant operations (e.g., a set of system calls relevant to the desired property). Then, the user finds all the sequences of these operations that violate the property, and encodes them using an FSA. Meanwhile, any execution of a program defines a *trace*, the sequence of security-relevant operations performed during that execution. MOPS uses the FSA to monitor program execution: as the program executes a security-relevant operation, the FSA transitions to a new state. If the FSA enters an *error state*, the program violates the security property, and this execution defines an *error trace*.

At its core, MOPS determines whether a program contains any feasible traces (according to the program's source code) that violate a security property (according to the FSA). Since this question is generally undecidable, MOPS errs on the conservative side: MOPS will catch all the bugs for this property (in other words, it is sound, subject to certain requirements [6]), but it might also report spurious warnings. This requires the user to determine manually whether each error trace reported by MOPS represents a real security hole.

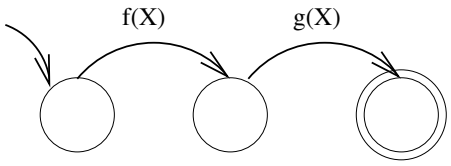

Figure 1. An FSA illustrating the use of pattern variables.

Specification of Security Properties. MOPS provides a custom property language for specifying security properties. The MOPS user describes each security-relevant operation using a syntactic pattern similar to a program's abstract syntax tree (AST). With wildcards, these patterns can describe fairly general or complex syntactic expressions in the program. The user then labels each FSA transition using a pattern: if the pattern matches an AST in the program during model checking, the FSA takes this transition.

To extend the expressiveness of these patterns, we introduced *pattern variables*, which can describe repeated occurrences of the same syntactic expression. For instance, if X denotes an pattern variable, the pattern $f(X, X)$ matches any call to the function f with two syntactically identical arguments. In any error trace accepted by an FSA, the pattern variable X has a single, consistent instantiation throughout the trace.

Formally, let Σ denote the set of ASTs. We may view a program trace as a string in Σ^*, and a property B as a regular language on Σ^*.[2] Pattern variables provide existential quantification over the set of ASTs. For instance, the pattern $\exists X.f(X, X)$ matches any call to f whose two arguments, once parsed, yield the same syntax subtree. If $B(X)$ is a language with an unbound pattern variable X, the language $\exists X.B(X)$ accepts any trace $t \in \Sigma^*$ where there exists an AST A' so that $B(A')$ accepts t. In other words, if $L(B)$ denotes the set of error traces accepted by the language B, we define $L(\exists X.B(X)) = \cup_{A'} L(B(A'))$. We use the convention that unbound pattern variables are implicitly existentially quantified at the outermost scope.

For instance, the FSA shown in Figure 1 accepts any trace containing a call to f followed by a call to g with the same syntactic argument. This property would accept traces such as $\langle f(0), g(0)\rangle$, $\langle f(a), g(a)\rangle$, or $\langle f(0), f(1), h(3), g(1)\rangle$ as error traces, but not $\langle f(0), g(1)\rangle$. See Appendix A for the algorithm we use to implement pattern variables.

Scalability. Since we aim at analyzing hundreds of large, real application, MOPS must be scalable in several senses. First, MOPS must run quickly on large programs. Second,

[2] Because the property can be described by an FSA, it is a regular language.

Property	*Reported Warnings*	*Real Bugs*	*Section*
TOCTTOU	790	41	3.1
Standard File Descriptors	56	22	3.2
Temporary Files	108	34	3.3
`strncpy`	1378	11*	3.4
Chroot Jails	1	0	Appendix C
Format String	(too many)	(unknown)	Appendix D
Total	2333	108	

Table 1. Overview of Results.

MOPS must run on different application packages without requiring the user to tweak each package individually.

MOPS achieved good performance through the use of *compaction* [6]. The basic idea is that, for any given property, most of the operations in the program are irrelevant to the property. By checking only those relevant operations in the program, MOPS runs much faster and requires much less memory. Compaction is extremely useful in keeping MOPS's runtime within reasonable limits in our experiments.

We have put much effort into integrating MOPS with existing build processes, including `make`, `rpmbuild`, and others. By interposing on `gcc`, the model checker sees the same code that the compiler sees. As a result, running MOPS on numerous software packages is as easy as invoking a MOPS script with the names of these packages. This ease of use has been critical to our success in checking such a large number of packages.

Error Reporting. MOPS reports potential errors in a program using error traces. A typical problem with reporting error traces is that a single bug can cause many (sometimes infinitely many) error traces. To avoid overloading the user, MOPS divides error traces into groups such that each group contains all the traces caused by the same bug. More precisely, two traces belong to the same group if the same line of code in both traces causes the FSA to enter an error state for the first time via the same transition[3]. The user can then examine a representative trace from each group to determine whether this is a bug and, if so, to identify the cause of the bug.

Not all error traces identify real bugs: imprecision in the analysis causes spurious traces. MOPS provides an HTML-based user interface (Appendix B) where the user can examine traces very rapidly. The user, however, does spend time identifying false positives, so the cost of using MOPS correlates roughly to the number of trace groups, each of which the user has to examine. In our experiments, we quantify the cost of using MOPS by measuring the number of false positives, counting only one per trace group.

[3]This implies that both traces enter the same error state. An FSA may contain multiple error states, corresponding to different kinds of bugs.

Resource Usage. The running time of the model checker is usually dwarfed by the time a human spends perusing error traces. Still, since our goal is to audit entire distributions, we have aimed to make computation time small. We timed the process of model checking several of our properties against all of Red Hat 9. Using MOPS to look for TOCTTOU vulnerabilities (filesystem races) among all Red Hat 9 packages requires about 1 GB of memory and takes a total of 465 minutes—a little less than 8 hours—on a 1.5 GHz Intel Pentium 4 machine. Detecting temporary file bugs takes 340 minutes of CPU time and about the same memory footprint. The observed wall-clock time was between 20% and 40% more than the CPU time. MOPS produces an extraordinary amount of output, and is required to read in extremely large control flow graphs; I/O thus constitutes a significant portion of this running time, although it is dominated by the time needed for model checking itself.

Also of chief concern to us was being able to audit manually all error traces produced by MOPS. Error trace grouping was a huge time saver: a typical group has more than 4 traces, but some groups contain more than 100 traces. The amount of human effort that was spent auditing the error groups is roughly proportional to the total number of groups. We spent about 100 person-hours auditing error reports from the TOCTTOU property, 50 person-hours for the temporary file property, and less for the other properties. Several of us were undergraduate students who had no prior experience with MOPS prior to joining this project.

3 Checking Security Properties

We developed six security properties. Table 1 shows a summary of the bugs discovered. For each property, the table shows the number of warnings reported by MOPS, the number of real bugs, and the section that describes detailed findings on this property. We will describe four properties in detail, explain the model checking results, and show representative bugs and vulnerabilities that we discovered in this section. We will defer the discussion on the last two properties, on which we did not find real bugs, to the appendix.

Red Hat Linux 9 consists of 839 packages with 60.0

million lines of code, the majority of which are integer-style programs written in C. They can all be built using the `rpmbuild` tool. The largest of the packages is around 4667K lines of code, and the smallest is less than a thousand.

3.1 TOCTTOU

Race condition attacks have perennially plagued secure systems. One common mistake is that a program checks the access permission of an object and, if the check succeeds, makes a privileged system call on the object. For example, one notorious error involves the `access()` and `open()` system calls. Consider a program running as root (e.g., setuid root, or executed by root) executing the following code fragment:

```
if (access(pathname, R_OK) == 0)
  fd = open(pathname, O_RDONLY);
```

The programmer is attempting to enforce a stricter security policy than the operating system. However, the kernel does not execute this sequence of system calls atomically— so if there is a context switch between the two calls or if the program is running on a multiprocessor system, another program may change the permission of the object. When the above program resumes its execution, it then blindly performs `open()` even though the user should no longer have access permission to the object.

Another example comes from UNIX folklore. It is well known that the root user should not recursively remove files inside directories that may be writable by other users. For example, "`rm -rf /tmp`" is a dangerous command, even if root has verified that the directory `/tmp` contains no symlink to other parts of the file system. The reason is that after `rm` verifies that a directory is not a symlink but before it enters the directory to delete the files within, an adversary may replace the directory with a symlink to another part of the file system, therefore tricking `rm` into deleting that part of the file system.

Many of the vulnerabilities that we found are exploitable when two users share access to some portion of the file system, and one user is operating on the shared portion while the other mounts an attack by replacing symbolic links. Although programs commonly attempt to ensure that they do not follow symbolic links before doing dangerous operations like `unlink()`, they often check it incorrectly and are susceptible to TOCTTOU attacks as a result.

We have experimented with several different FSAs to capture these types of vulnerabilities. In our first attempt, we chose an FSA that had two transitions—one from the start state to an intermediate state, and the other from the intermediate state to the accepting state. Both transitions

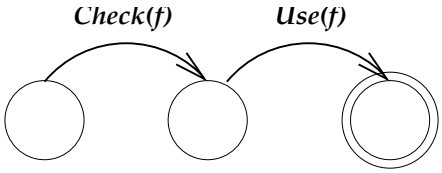

Check(f) = {`stat(f), lstat(f), access(f), readlink(f), statfs(f)`}.
Use(f) = {`basename(f), bindtextdomain(f), catopen(f), chown(f), dirname(f), dlopen(f), freopen(f), ftw(f), mkfifo(f), nftw(f), opendir(f), pathconf(f), realpath(f), setmntent(f), utmpname(f), chdir(f), chmod(f), chroot(f), creat(f), execv(f), execve(f), execl(f), execlp(f), execvp(f), execle(f), lchown(f), mkdir(f), fopen(f), remove(f), tempnam(f), mknod(f), open(f), quotactl(f), rmdir(f), truncate(f), umount(f), unlink(f), uselib(f), utime(f), utimes(f), link(f), mount(f), rename(f), symlink(f)`}.
Non-filename arguments are omitted.

Figure 2. A refined FSA for finding TOCTTOU (file system race condition) vulnerabilities.

were defined on the union of the file system calls that access the file system using a pathname argument. However, we found that this naive property results in too many false positives—for example, `chdir(".")` followed by a `chdir(".")` would trigger a false positive. Typically these situations arise when a single system call is located inside of a loop, and both transitions in the FSA are made as a result of executing the same line of code. Since these are obviously not security holes, we decided to separate out the file system calls that can be classified as "checks" from those that are "uses".

We refined the property by dividing the file system calls into two distinct sets. Figure 2 shows the general structure of the FSA. We assume here, and in subsequent illustrations, that there is an implicit *other* transition from every state back to itself; if none of the normal outgoing transitions match, the *other* transition is taken, and the FSA stays in the same state. The intuition is as follows: a call to *Check(f)* is probably intended to establish an invariant (e.g., "*f* is not a symlink"), and a call to *Use(f)* might rely on this invariant. Of course, these invariants might be violated in an attack, so *Check(f)* followed by *Use(f)* may indicate a TOCTTOU vulnerability. This refined property is much more manageable and finds most of the bugs we are interested in. However, the more general property is capable of finding some bugs which the narrower TOCTTOU cannot—for example, `creat(f)` followed by `chmod(f)`.

The types of vulnerabilities we have found can be classified under the following categories:

1. [*Access Checks*] A check is performed—often by a program running as root—on file access permissions. The result of the check is then used to determine whether a resource can be used. The `access(f)` and `open(f)` race at the beginning of this section illustrates this class of bugs.

2. [*Ownership Stealing*] A program may `stat()` a file to make sure it does not exist, then `open()` the file for writing. If the `O_EXCL` flag is not used, an attacker may create the file after the `stat()` is performed, and the program will then write to a file owned by the attacker. We consider this a vulnerability, because the program may disclose sensitive information.

3. [*Symlinks*] Vulnerabilities due to symbolic links arise when two users share the same portion of the file system. One user can change a file to a symlink to trick the other user to mistakenly operate on an unexpected file. The method of such an attack depends on whether the system call follows symlinks. Broadly, there are two classes of system calls:

   (a) [*Syscalls that follow symlinks*] Many system calls will follow symbolic links that occur anywhere in the pathname passed to them. These present no barrier to attack.

   (b) [*Syscalls that don't follow symlinks*] Other system calls avoid following symbolic links if they occur in the last component of their pathname argument. For instance, if `c` is a symbolic link to `d`, calling `unlink("/a/b/c")` will delete the symbolic link itself rather than the target of the link: it deletes `/a/b/c`, not `/a/b/d`. However, many programmers do not realize that these calls will gladly follow any symlinks that occur in earlier components of the pathname. For example, if `b` is a symlink to `../etc`, then `unlink("/a/b/passwd")` will delete the password file `/etc/passwd`. Consequently, to attack this second class of system calls, it suffices for the attacker to tamper with one of the intermediate components of the path.

Table 2 in Appendix E shows all the TOCTTOU bugs we discovered. The third column shows the number of traces that MOPS thought violated the security property, and the last column shows the number of those traces that were actual bugs. We successfully mounted attacks on several programs to verify that the bugs were real. Many bugs we found were not previously known. Some were previously reported (but apparently not yet fixed and not known to us at the time of our experiments). To illustrate the kinds of bugs we found with MOPS, we will show three representative examples of TOCTTOU bugs.

```
# binutils-2.13.90.0.18-9 :: ar #
exists = lstat (to, &s) == 0;
/* Use rename only if TO is not a symbolic
   link and has only one hard link.  */
if (! exists || (!S_ISLNK (s.st_mode)
  && s.st_nlink == 1)) {
  ret = rename (from, to);
  if (ret == 0) {
    if (exists) {
      chmod (to, s.st_mode & 0777);
      if (chown (to, s.st_uid,
           s.st_gid) >= 0) {
        chmod (to, s.st_mode & 07777);
      }
  ...
```

In our first example, the program *ar* executes the code fragment above to replace an archive file with one of the same name. It calls `lstat` on the destination file and then checks if the file is a symbolic link. If it is not, *ar* calls `rename` on the file and then sets the mode of the file. This code, however, is unsafe. An adversary may change the file to be a symbolic link after *ar* checks for symbolic links. Then, *ar* will happily change the mode of whatever the symbolic link points to—assuming the user running *ar* has permission to do so. The attack is applicable when two users have write access to the directory of the archive file.

```
# initscripts-7.14-1 :: minilogd #
/* Get stat info on /dev/log so we can later
   check to make sure we still own it... */
if (stat(_PATH_LOG,&s1) != 0) {
   memset(&s1, '\0', sizeof(struct stat));
}
...
if ( (stat(_PATH_LOG,&s2)!=0) ||
     (s1.st_ino != s2.st_ino ) ||
     (s1.st_ctime != s2.st_ctime) ||
     (s1.st_mtime != s2.st_mtime) ||
     (s1.st_atime != s2.st_atime) ) {
         done = 1;
         we_own_log = 0;
}
/* If we own the log, unlink it before trying
   to free our buffer. Otherwise, sending the
   buffer to /dev/log doesn't make much sense */
if (we_own_log) {
   perror("wol");
   unlink(_PATH_LOG);
}
```

The second code fragment is taken from the program *minilogd*, which is run by root. This program may unlink `_PATH_LOG` (which is defined to be `/dev/log` by default) if

it thinks it exclusively owns the file. It compares the timestamps on the file at two different times in the execution of the program and, if they are equal, decides that it exclusively owns the file and then removes the file. However, even if another program modifies the file after *minilogd* checks the timestampes, *minilogd* will still unlink it, possibly corrupting other programs. An additional vulnerability exists when user programs can write to the log file; for instance, if `_PATH_LOG` is defined as something like `/home/alice/log` instead. In this case, Alice can trick *minilogd* into removing anything on the file system. We have found many vulnerabilities that are similar to the latter case. It is important that these filename constants be checked very carefully when these programs are built, since it may not be obvious to users that defining `_PATH_LOG` to a user-writable file can result in a total compromise of the file system.

```
# zip-2.3-16 :: zip #
  d_exists = (LSTAT(d, &t) == 0);
  if (d_exists) {
    /* respect existing soft and hard links! */
    if (t.st_nlink>1 ||
          (t.st_mode&S_IFMT)==S_IFLNK)
       copy = 1;
    else if (unlink(d))
       return ZE_CREAT;
  }
```

The final code snippet comes from the widely-used program *zip*. If the destination file already exists, *zip* will move the file to a new location, unlink the old copy, and write the new copy. The program verifies that the file is not a link before calling `unlink` on it. The attack is applicable when two users share a portion of the file system and one user is running *zip* to write a new file to the shared space. If the other user is malicious, after *zip* calls `stat`, the user can change the file to be a symbolic link that points to another part of the file system. Since `unlink` will not follow the last component of a pathname, the attacker would have to change one of the components in the middle of the pathname to a symbolic link. For instance, if Alice is using *zip* to write a file to `/shared/alice/afile`[4], Bob can change `/shared/alice` to be a symbolic link that points to `/home/alice`. Then the *zip* program running on behalf of Alice will remove `/home/alice/afile`. Most users will not be aware that using a shared directory enables such attacks, so it seems unfair to blame Alice for doing so. In this case, *zip* does try to do the right thing by checking for symbolic links; it just happens to get the check wrong.

[4]The suggested scenario requires that the sticky bit is not set. The sticky bit prevents deletion of files and directories for anyone except the creator, even if others have write access. Generally, `/tmp` has the sticky bit set.

3.2 A Standard File Descriptor Attack

The first three file descriptors of each Unix process are called *standard file descriptors*: FD 0 for the standard input (*stdin*), FD 1 for the standard output (*stdout*), and FD 2 for the standard error (*stderr*). Several commonly used C standard library functions read from or write to these standard file descriptors; e.g., `fgets()` reads from *stdin*, `printf()` writes to *stdout*, and `perror()` writes to *stderr*. Programs that print information intended for the user to see, or diagnostic information, typically do so on FDs 1 and 2. Customarily, a program starts with its standard file descriptors opened to terminal devices. However, since the kernel does not enforce this convention, an attacker can force a standard file descriptor of a victim program to be opened to a sensitive file, so that he may discover confidential information from the sensitive file or modify the sensitive file.

For example, suppose a victim program is setuid-root[5] and executes the following code:

```
/* victim.c */
fd = open("/etc/passwd", O_RDWR);
if (!process_ok(argv[0])) perror(argv[0]);
```

Then the adversary can run the following attack program to exploit the standard file descriptor vulnerability in the victim program:

```
/* attack.c */
int main(void) {
  close(2);
  execl("victim",
    "foo:<pw>:0:1:Super-User-2:...", NULL);
}
```

This attack works as follows. First, the attack program closes FD 2 and executes victim.c. A child process will inherit the file descriptors from the parent process; consequently, the victim program starts with FD 2 closed. Then, when the victim opens the password file */etc/passwd*, the file is opened to the smallest available file descriptor—in this case, FD 2. Later, when the victim program writes an error message by calling `perror()`, which writes to FD 2, the error message is appended to */etc/passwd*. Due to the way the attacker has chosen this error message, the attacker may now log in with superuser privileges. These bugs are particularly dangerous when the attacker can influence the data written to the standard FD.

In the previously discussed vulnerability, the attacker is able to append content to an important system file. One can envision similar attacks on the *stdin* file descriptor in which the attacker can read content from a file that is not intended

[5]A setuid-root program runs with root privileges, even if it is executed by an unprivileged user.

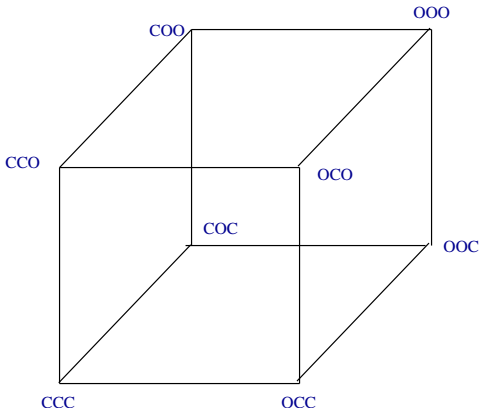

Figure 3. The structure of an FSA for finding file descriptor vulnerabilities. This FSA tracks the state of the three standard file descriptors across `open()` calls.

to be publically available. To stage such an attack, the malicious program would first close FD 0, then execute the privileged program containing code that unwittingly tries to read from *stdin*. The vulnerable program will now instead read data that is from the attacker's choosing, and possibly disclose confidential information in the process. Note that the program must relay the information it has learned back to the attacker, either directly or through a covert channel. The latter means of disclosure is impossible to detect using the single program analysis techniques we employ.

The way to prevent these types of attacks is simple—a program that runs as setuid-root should ensure that the three lowest numbered file descriptors are opened to a known safe state prior to carrying out any important operations. A common way to do this is by opening a safe file, such as */dev/null*, three times when the program first starts up. In the case that someone tries to attack the program by closing one or more of the file descriptors prior to executing the victim program, no harm is done because they are re-opened to point to */dev/null*. This solution is usually acceptable because the overhead is only three system calls. In the case that all 3 FDs were already opened, the program also consumes three file descriptor slots.

We have developed two MOPS properties intended to detect this types of attack. A description of one of the automata is described below; the second is very similar, and we later explain how it is formed.

Our FSA used in this property (see Figure 3) contains eight states that are used to describe a unique combination of the states of the three standard file descriptors 0, 1 and 2. For example, the state `OCC` represents that FD 0 is open, but FD 1 and FD 2 are closed. The program may start in any of the seven states where at least one of the three standard FDs is closed; the case where all of the standard FDs are initially open is the usual one, and not of interest to an attacker. The starting state will be chosen nondeterministically during the model checking phase to insure all possibilities are explored. Transitions in the FSA occur only along the edges of the cube, as there are no system calls that can change the status of multiple standard file descriptors at once.

The basic FSA structure in Figure 3 is not entirely complete, as we have not shown the error state. For detecting the class of attacks which can cause the vulnerable program to write to arbitrary files, we add a new error state and a transition to the error state when a file that is neither */dev/null* nor */dev/tty* is opened on FD 1 or 2 in a mode other than read-only. For detecting the class of attacks that may disclose the contents of secret files, we add transitions to the error state from the four states in which FD 0 is closed (`COO`, `COC`, `CCO`, and `CCC`), and a file other than */dev/null* or */dev/tty* is opened for reading. These two transitions are seperated into two different automata, to give the two properties.

To save space, we have not labeled the transitions along the edges of the cube. These transitions are taken for system calls that are considered a "safe" open—that is, when */dev/null* or */dev/tty* is opened. For example, if the current state is `COC` and a "safe" open is encountered, then the new state is `OOC`, since the file will be opened on the lowest-numbered available FD.

We have audited the programs that run as setuid root on our Linux distribution, and have identified a number of bugs (but not exploitable vulnerabilities at this time). In many cases, an attacker can cause a setuid program to write data not of her choosing to temporary files, lock files, or PID files (files used to store the process ID of the currently running service). These situations can be potential vulnerabilities if some other program trusts the contents of the PID file. For example, consider a system administration script for restarting some network daemon that executes `kill ‘cat pidfile‘`. If the attacker exploits a setuid program that writes to this PID file to introduce a line of the form "*PID*`; rm /etc/passwd`" into the PID file, then the administration script might unwittingly remove `/etc/passwd` when it is next run. We have not yet found any fully exploitable scenario like this, but the fact that some setuid programs allow corrupting PID files like this is perhaps room for some concern.

An example of a bug we found in the program *gnuchess*, a chess playing application, follows:

```
int main(int argv, char *argv[]){
   ...  BookBuilder(depth, ...); ...
}
void BookBuilder(short depth, ...){
  FILE *wfp,*rfp;
  if (depth == -1 && score == -1) {
    if ((rfp = fopen(BOOKRUN,"r+b")) != NULL) {
      printf("Opened existing book!\n");
```

```
    } else {
      printf("Created new book!\n");
      wfp = fopen(BOOKRUN,"w+b");
      fclose(wfp);
      if ((rfp = fopen(BOOKRUN,"r+b"))
            == NULL) {
        printf("Could not create %s file\n",
            BOOKRUN);
       return;
      }
      ...
}
```

The function `BookBuilder` is called to manipulate and read from the playbook used by the game. Although there is no attack to compromise security, it is easy to see the bug. The playbook can become corrupted when a malicious user closes all file descriptors except standard out, and invokes the gnuchess program. The file *BOOKRUN* will then be opened onto standard out, and the subsequent writes from *printf()* can corrupt the book.

Table 3 shows the setuid applications for which we found bugs, and the results from running MOPS and our two file descriptor properties on them. There were two main sources of false positives: 1) the property does not track the UID privilege changes inside the program, so the program may drop privileges before opening files, and 2) the property did not recognize that the program safely opened */dev/null* three times, due to a nonstandard invocation of safe opens. Unfortunately these are difficult false positives to recognize, because they require the user to look at the trace in its entirety as opposed to the usual points of interest (line numbers that caused transitions in the FSA). The presentation of our results differentiates betweens bugs and exploits. For this property, we classify bugs as programming mistakes that can cause unexpectd program behavior, but not necessarily lead to any compromise of security. For example, an attack that can compromise the contents of a non-important file, such as a lockfile, falls under the category of a bug. An exploit needs to have security concerns—we have found none of these to date. However, it was surprising that many setuid programs did not open */dev/null* three times before performing file operations, given that it has low overhead and guarantees safety with regards to this property.

3.3 Secure Temporary Files

Applications often use temporary files as a means for passing data to another application, writing log information, or storing temporary data. Often times on a Unix system, the files will be created in the `/tmp` directory, which is world writable and readable. For example, the GNU C compiler creates temporary files when it is compiling programs, and later passes them to the linker. Many of the functions to create temporary files that are found in the C standard library are insecure. The reason is that they do not return a file descriptor, but rather a file name. An adversary that is able to predict the filename can thus create the file before the application has a chance to open or create it. This attack can give the adversary ownership of the temporary file, which is undesirable[6].

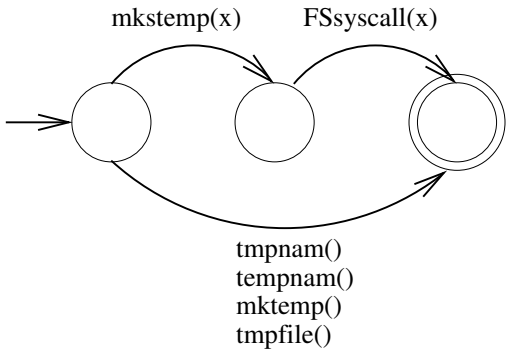

Figure 4. An FSA to detect insecure uses of temporary files.

We identified the set of insecure functions: `mktemp`, `tmpnam`, `tempnam`, and `tmpfile`. These functions should never be used. There is one function that can be secure, depending on how it is used: `mkstemp`. Security requires that the filename retrieved from `mkstemp` is never subsequently used in another system call: `mkstemp` returns both a file descriptor and a filename, but a secure program should not use the filename. Figure 4 illustrates our automaton.

In some situations a program may use an insecure function but still not have a vulnerability. An example of such a situation is a program that uses `tempnam` to create a temporary filename in `/tmp`, then attempts to open that file using the UNIX `O_CREAT` and `O_EXCL` flags. Combined, the flags ensure that the file does not already exist. The program can then check the return value from `open` to verify that an attacker did not first create the file. It is possible to modify the simple FSA presented in Figure 4 to detect these situations and not report them as false positives, however it makes the FSA much larger. In practice, it seems prudent to manually examine these situations anyway. One can imagine examples in which a program uses the flags but does not take the correct action based on the return value from `open`.

Table 4 shows the applications that violated the property by using one of the insecure calls. Table 5 shows the subset of bugs due to inappropriate use of the filename returned by `mkstemp`. These tables may be found in Appendix E.

Below we show a representative example of a program that violates the clause of our property that finds reuses

[6] In the `gcc` example, an adversary could insert malicious code into a user's program by replacing the temporary file with the desired code.

of the parameter passed to mkstemp. Not only is this the most complicated example presented thus far, but it shows how the whole-program inter-procedural analysis was effective. The code will be presented in several fragments as they occur temporally while executing the program. The example comes from the program *yacc* from the package byacc-1.9-25.

```
static void open_files() {
  int fd;
  create_file_names();
  if (input_file == 0) {
    input_file = fopen(input_file_name, "r");
    if (input_file == 0)
       open_error(input_file_name);
    fd = mkstemp(action_file_name);
    if (fd < 0 || (action_file =
      fdopen(fd, "w")) == NULL) {
      if (fd >= 0)
        close(fd);
      open_error(action_file_name);
    }
}
```

Before the above program fragment executes, there is some setup code that sets the value of the variable action_file_name. Specifically, it is a string whose first component is a pathname to a temporary directory (by default, it chooses /tmp, but this can be changed by defining an environment variable), and whose second component is a temporary file template[7]. The above code alone should be of concern to us. Recall that mkstemp returns a file descriptor that can be safely used, but the template passed to mkstemp is not safe to re-use. In this case, we see the template being passed to another function called open_error:

```
void open_error(char *filename) {
  warnx("f - cannot open \"%s\"", filename);
  done(2);
}
```

From the above fragment it looks like the function warnx may be a good candidate for inspection because it is the recipient of the filename we are interested in tracking. Strangely enough, MOPS directs us to the function done:

```
void done(int k) {
  if (action_file)
    fclose(action_file);
  if (action_file_name[0])
    unlink(action_file_name);
```

[7] A template is a partial filename with a number of placeholders denoted by a special character X that will be filled in with random numbers by the function creating the temporary filename.

Here we find the bug. The variable action_file_name, which is the template passed to mkstemp, is re-used as an argument to the unlink system call. This is unsafe. By the time we call unlink, the filename may no longer point to the location we think it does. Recall that the directory in which the file is being created may be world writable. An attacker that has write access to the file can change it to a symbolic link, and cause the program to unlink other unexpected files on the system. Unfortunately there does not appear to be a good resolution to the problem.

In addition to these two types of vulnerabilities, we have also investigated a third class that can arise on systems with old versions of *glibc*. In previous versions, the library functions for creating temporary files automatically set the permissions on the file to be world writable and readable by default. Therefore, programs must use the umask system call to change default file permissions before creating temporary files; if they do not, they will be vulnerable to attack. We used MOPS to find programs that use temporary files and do not first issue a umask call. Formally, if there exists a path from the program entry to a function call that creates a temporary file, and no umask call exists on the path, then MOPS will report an error.

We found that the overwhelming majority (greater than 90%) of programs do not use umask and thus would be vulnerable if used with old versions of *glibc*. However, the severity of these bugs is very low, and they are large in number, so we have not enumerated them here.

3.4 Attacks Against strncpy

There are several common attacks against programs that misuse the standard library function strncpy. strncpy(d,s,n) copies a string of characters pointed to by s into the memory region designated by d. If s contains more than n characters, strncpy only copies the first n characters. If s contains fewer than n characters, strncpy copies all the characters in *s* and then fills *d* with null characters until the length of d reaches n.

strncpy is easy to misuse for two reasons. First, it encourages off-by-one errors if the programmer is not careful to compute the value of n precisely. Off-by-one errors can often cause the program to write past the end of an array bounds, which can in turn lead to buffer overrun attacks against the program. In particular, consider a case where the string buffer in question is allocated on the runtime stack (as it will be when the buffer is an array local to a function in C), and the user of the program is able to control the contents of the source of s. If the program writes past the end of the buffer, a malicious user may construct a special string s, such that when the program writes past the array bounds, it writes special code into the stack frame that corrupts the program. Secondly, because the function does not automat-

ically null-terminate a string in all cases (for instance, when the size of the source string is larger than n), it is a common mistake for a program to create unterminated strings during its execution.

We have constructed an FSA to try and catch both scenarios as described above (see Figure 8). The intuition is that we identify several idioms that are correct ways to null terminate a string, and raise an alarm when one of these idioms is not used. For example, a common idiom is the following code sequence that is safe:

```
buf[sizeof(buf)-1] = '\0';
strncpy(buf, ..., sizeof(buf)-1);
```

In the above case, the buffer will also be terminated. However, the following two cases show a common misuse of strncpy:

- `buf[sizeof(buf)-1] = '\0';`
  `strncpy(buf, ..., sizeof(buf));`
- `memset(buf, 0, sizeof(buf)-1);`
  `strncpy(buf, ..., sizeof(buf)-1);`

In the first unsafe example, the string is null-terminated before the strncpy, and the execution of the function subsequently may overwrite the null-terminating character. In the second unsafe example, memset is used to zero-out the destination buffer; unfortunately, it is misused—the third argument needs to be the size of the entire buffer. Our FSA attempts to detect patterns that appear fishy, and alerts the user to their presence. Pattern variables are used judiciously to make MOPS precisely match the null-terminating code to the strncpy code that uses the same buffers. The property requires more manual inspection than other properties, because we have chosen an approach where we identify correct behavior, then raise an alarm at code that does not match our expectations. Moreover, our property does not attempt to find all strncpy bugs, focusing instead on patterns that are particularly suspicious.

Our strncpy FSA has alerted us to a number of bugs. Below we show one of the most interesting examples. It comes from the program *xloadimage*:

```
void dumpImage(Image *image, char *type,
  char *filename, int verbose) {
  int a;
  char typename[32];
  char *optptr;
  optptr = index(type, ',');
  if (optptr) {
    strncpy(typename, type, optptr - type);
    typename[optptr - type] = '\0';
  ...
```

In the above code fragment, MOPS identifies an idiom that does not appear to be safe. The character buffer typename is declared to be 32 bytes long, but the length passed to strncpy is computed entirely based on the second argument to the function dumpImage. We must verify that this string cannot be constructed in such a way that when optptr - type is computed, the result is longer than 32 bytes. MOPS is able to direct us to the call site of this function, in which we see the following (abbreviated):

```
newopt->info.dump.type= argv[++a];
...
dumpImage(dispimage, dump->info.dump.type,
  dump->info.dump.file, verbose);
```

Shockingly, the contents of the string come from the command line arguments, and can be set entirely by the user. Consequently, a malicious user can supply a carefully crafted argument to the function which causes the function dumpImage to write past the end of an array, causing a buffer overrun.

Unfortunately, this property produced 1378 unique warnings, too many to examine exhaustively by hand. Therefore, we picked a semi-random sample of 16 packages out of the set of 197 packages with one or more warning, yielding a set of 53 warnings. Examining these 53 warnings revealed 11 bugs spread among 6 packages, where in each case a string could be left unterminated or the strncpy() operation could overflow buffer bounds for some input. We did not attempt to assess the security implications of these bugs, though we expect that many of them could be exploited under some circumstances.

Based on this limited sample, we suspect that a full manual audit using MOPS would turn up many more strncpy() bugs. Since we saw 11 bugs among 53 warnings, this suggests a false positive rate of about 79%, and a true positive rate of about 21%. If all warnings are equally likely to correspond to real bugs, we might estimate that there are about $1378 \times \frac{11}{53} \approx 286$ bugs, with a 95% confidence interval of between 165 to 468 bugs. Of course, these estimates are fairly rough, but our best prediction is that a full manual audit of all the MOPS warnings would turn up in excess of one hundred strncpy() bugs.

4 Related Work

There is a broad and growing array of work on software model checking, and MOPS represents just one of several tools in this area. BLAST [12] and SLAM [2] are dataflow-sensitive model checkers that use adaptive iterative refinement to narrow down the locations of bugs. Both have been used primarily on smaller programs, such as device drivers, but they are able to provide a much more precise analy-

sis. Similar to BLAST is MAGIC [4], a system that abstracts predicates and uses a theorem prover for dataflow analysis. CMC, a model checker for C and C++ programs [14], has been used on large-scale applications like the entire Linux kernel [11]. Metal, a ground-breaking bugfinding tool which is similar in concept to a model checker, has been very successful at finding a broad variety of rule violations in operating systems [1, 10]. Metal has been augmented with Z-ranking, a powerful technique for reducing the number of false positives, and used to find many bugs in large applications [13]. MOPS has previously been used to find bugs in eight security-relevant packages from the Red Hat distribution [5, 6]. However, none of these tools have yet been applied on as large a scale as shown in this paper.

File system race conditions have been extensively studied in the computer security literature. Bishop and Dilger first articulated the vulnerability pattern and developed a syntactic pattern-matching analysis for detecting TOCTTOU vulnerabilities in C code [3]; however, because their analysis is not semantically based, it is unable to find many of the vulnerabilities found in this work. Dean and Hu propose a probabilistic defense that makes it very difficult for adversaries to win races [8]; however, their approach requires re-writing existing code. Also, some authors have proposed runtime program analysis methods to detect TOCTTOU bugs by monitoring program executions and preventing their exploitation [7]; our work differs by trying to find TOCTTOU bugs at compile time, rather than at runtime.

Today, the state of the art for most practitioners involves tools that scan program source code for simple syntactic patterns [15, 20]. For instance, RATS [16], ITS4 [18], Flawfinder [19], and PScan [9] are all based on analyzing the stream of tokens produced by the lexer to look for things like function calls to certain error-prone library calls. MOPS is able to look more deeply, both because of its ability to match on complex AST patterns and because of its ability to look for certain sequences of operations by using control-flow analysis. As a result, all of the properties described in this paper go beyond what those lexical tools can find. Nonetheless, the popularity of even very simple tools demonstrates a growing interest in security auditing tools.

5 Conclusion

Our work demonstrates that large-scale model checking is feasible. We showed that it is possible to develop models of incorrect and insecure program behavior that are precise enough to prevent false positives from dwarfing the real bugs; also, as this work showed, many of these properties can be encoded in MOPS without serious loss of soundness. Thanks to the sophisticated error reporting in MOPS, we found that we were able to manually inspect all error traces. Consequently, we were able to find many (108) real exploitable software bugs; in several cases, we have crafted attacks to verify their validity. As a result of this experience, we are convinced that software model checking can easily be integrated into the development process, particularly when using model checkers like MOPS that can be integrated into build processes at the highest level.

References

[1] K. Ashcraft and D. Engler. Using programmer-written compiler extensions to catch security holes. In *Proceedings of IEEE Security and Privacy 2002*, 2002.

[2] T. Ball and S. K. Rajamani. The SLAM project: Debugging system software via static analysis. In *POPL '02: Proceedings of the ACM SIGPLAN-SIGACT Conference on Principles of Programming Languages*, 2002.

[3] M. Bishop and M. Dilger. Checking for race conditions in file accesses. *Computing Systems*, 9(2):131–152, Spring 1996.

[4] S. Chaki, E. Clarke, A. Groce, S. Jha, and H. Veith. Modular verification of software components in C. In *International Conference on Software Engineering*, pages 385–395, May 2003.

[5] H. Chen, D. Dean, and D. Wagner. Model checking one million lines of C code. In *Proceedings of the 11th Annual Network and Distributed System Security Symposium*, San Diego, CA, Feb. 4–6, 2004.

[6] H. Chen and D. Wagner. MOPS: an infrastructure for examining security properties of software. In *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS)*, pages 235–244, Washington, DC, Nov. 18–22, 2002.

[7] C. Cowan, S. Beattie, C. Wright, and G. Kroah-Hartman. RaceGuard: Kernel protection from temporary file race vulnerabilities. In *Proceedings of the Tenth USENIX Security Symposium*, 2001.

[8] D. Dean and A. Hu. Fixing races for fun and profit. In *Proceedings of the Thirteenth Usenix Security Symposium*, Boston, MA, 2004.

[9] A. DeKok. PScan: A limited problem scanner for C source files. Available at `http://www.striker.ottawa.on.ca/aland/pscan`.

[10] D. Engler, B. Chelf, A. Chou, and S. Hallem. Checking system rules using system-specific, programmer-written compiler extensions. In *OSDI*, 2000.

[11] D. Engler and M. Musuvathi. Model-checking large network protocol implementations. In *Proceedings of the First Symposium on Networked Systems Design and Implementation*, 2004.

[12] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Software verification with BLAST. In *Proceedings of the 10th SPIN Workshop on Model Checking Software*, 2003.

[13] T. Kremenek and D. Engler. Z-Ranking: Using Statistical Analysis to Counter the Impact of Static Analysis Approximations. In *Proceedings of the 2003 Static Analysis Symposium*, 2003.

[14] M. Musuvathi, D. Park, A. Chou, D. R. Engler, and D. L. Dill. CMC: A Pragmatic Approach to Model Checking Real Code. In *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation*, Dec. 2002.

[15] J. Nazario. Source code scanners for better code, Jan. 2002. Available at `http://www.linuxjournal.com//article.php?sid=5673`.

[16] Secure Software, Inc. RATS. Available at `http://www.securesoftware.com/download_rats.htm`.

[17] U. Shankar, K. Talwar, J. S. Foster, and D. Wagner. Detecting format string vulnerabilities with type qualifiers. In *Proceedings of the Eleventh Usenix Security Symposium*, pages 201–218, 2001.

[18] J. Viega, J. Bloch, T. Kohno, and G. McGraw. ITS4: A Static Vulnerability Scanner for C and C++ Code. In *16th Annual Computer Security Applications Conference*, Dec. 2000.

[19] D. Wheeler. Flawfinder. Available at `http://www.dwheeler.com/flawfinder/`.

[20] J. Wilander. A comparison of publicly available tools for static intrusion prevention. In *7th Nordic Workshop on Secure IT Systems*, Nov. 2002.

A Efficient Model Checking with Pattern Variables

A.1 Current implementation

In the current version of MOPS, pattern variables are implemented in a very simple way. In a first pass, for each AST pattern—say, $f(X, X)$—mentioned in the FSA and containing a pattern variable, MOPS finds all potential matching subtrees in the program's parsed source code (treating the pattern variable temporarily as a wildcard). This is used to derive a set of possible instantiations for the pattern variable X, and this process is performed separately for each pattern variable occurring in the security property. In the second pass, MOPS uses these sets to exhaustively enumerate all possible assignments to the set of pattern variables. For each assignment, the pattern variables are instantiated and the model checker is applied. If there are k pattern variables, each with n possible instantiations, the model checker is invoked n^k times.

Formally, this algorithm may be described as follows. Let $B(X_1, \ldots, X_k)$ be an FSA with pattern variables $X_1, \ldots, X_k$. Let P denote the program, expressed as a pushdown automaton, and let $\text{Stmts}(P)$ denote the set of statements (ASTs) in P. Let $L(P) \subseteq \Sigma^*$ denote the context-free language accepted by P, and similarly $L(B(\cdots)) \subseteq \Sigma^*$ the regular language accepted by $B(\cdots)$. Our algorithm is:

MATCHES$(B(X_1, \ldots, X_k), P)$:
1. For $i = 1, \ldots, k$, do:
2. Set $S_i := \emptyset$.
3. For each pattern $A(X_i)$ in B mentioning X_i, do:
4. If there is $s \in \text{Stmts}(P)$ matching $\exists X_i.A(X_i)$,
5. Let s' be the AST subtree such that $s = A(s')$.
6. Add s' to S_i.
7. Return $(S_1, \ldots, S_k)$.

CHECK$(B(X_1, \ldots, X_k), P)$:
1. Set $(S_1, \ldots, S_k) :=$ MATCHES$(B(X_1, \ldots, X_k), P)$.
2. For each k-tuple $(x_1, \ldots, x_k)$ with $x_i \in S_i$, do:
3. If $L(B(x_1, \ldots, x_k)) \cap L(P) \neq \emptyset$, output an error trace.

Line 3 of CHECK is implemented by a call to a pushdown model checker, and accounts for the overwhelming majority of the runtime of this scheme.

This simple algorithm allows us to leverage a core model checker with no built-in support for pattern variables, and it has worked fairly well for us as long as pattern variables are used sparingly. However, it does come with significant performance overhead when pattern variables are used heavily.

A.2 More Efficient Algorithms

We are aware of more efficient algorithms for model checking properties with pattern variables, but we have not yet implemented them. One optimization is to pre-screen each assignment using a flow-insensitive program analysis, like this:

FASTCHECK$(B(X_1, \ldots, X_k), P)$:
1. Set $(S_1, \ldots, S_k) :=$ MATCHES$(B(X_1, \ldots, X_k), P)$.
2. For each k-tuple $(x_1, \ldots, x_k)$ with $x_i \in S_i$, do:
3. If $L(B(x_1, \ldots, x_k)) \cap \text{Stmts}(P)^* \neq \emptyset$, then:
4. If $L(B(x_1, \ldots, x_k)) \cap L(P) \neq \emptyset$, then:
5. Output an error trace.

Note that Line 3 can be implemented efficiently, since it is a simple graph reachability problem: we mark each transition in B that matches some statement in P (when using the assignment $X_1 \mapsto x_1, \ldots, X_k \mapsto x_k$), then we use depth-first search to check whether B's error state is reachable from B's start state by following only marked transitions. Here we have used the fact that $L(P) \subseteq \text{Stmts}(P)^*$ and the observation that the intersection of two regular languages can be computed efficiently. This amounts to using a cheap flow-insensitive program analysis to quickly screen out assignments that can't produce any error traces, following up with a more expensive flow-sensitive, context-sensitive pushdown model checker only where necessary. Most assignments won't survive the cheap screening phase, so this optimization should improve performance significantly.

B User Interface for Error Reporting

Errors—i.e., violations of the security property—are reported to the user in HTML format. The user is shown a

brief explanation of what security rule was violated, accompanied by a set of error traces. Each error trace shows one path of execution through the program on which the property is violated, and thus explains why MOPS believes the program to be in error.

Our user interface (UI) displays error traces to the user concisely by highlighting FSA transitions and using an extended stack backtrace paradigm. Each statement in the trace where the property FSA changes states is likely to be particularly relevant to explaining the source of this error, so we always show each such transition. A full listing of all statements executed along this trace is usually far too long for a human to digest, so we summarize this information by showing a skeleton of the call stack. Any statement that does not cause an FSA transition is elided; any function whose body is entirely elided is itself elided; and this continues recursively. This leaves us with a tree in which each leaf is an FSA transition, and the path from the root to that leaf shows the stack backtrace at the time that FSA transition occurred.

An example may help clarify. Suppose we have a program where `main()` calls functions `f()`, `g()`, and `h()`, in that order; `f()` in turn calls `f0()` and then `f1()`; `g()` calls `g0()` and `g1()`; and similarly for `h()`. If the FSA transitions directly from the start state to the error state (in a single step) while executing the body of `g1()`, the error trace summary would look like this:

- `main()`
  - `g()`
    - `g1()`
      - Start → Error

Note that the calls to, and bodies of, `f()`, `f0()`, `f1()`, and `g0()` have been elided, because they had no effect on the state of the FSA. Alternatively, suppose that the FSA had changed state twice along this trace, once in `g1()` and then again later in `h0()`. In that case, a trace summary might look like this:

- `main()`
  - `g()`
    - `g1()`
      - Start → InProgress
  - `h()`
    - `h0()`
      - InProgress → Error

Trace summaries are interactive: clicking on one of the squares causes the body of that function to be expanded. This allows the user to see progressively more detailed error traces.

Our UI shows error traces in two panes, with a narrow left pane showing the trace summary and a right pane used to show the program source code. Clicking on any function call takes the right pane to the relevant call site in the source, and highlights that line of code. Clicking on an FSA transition goes to the statement that triggered the FSA to change states. Our experience is that if there is a bug in the program, it is most likely to be found where the last transition (the one that first enters the FSA's error state) occurs in some error trace. Therefore, by default our UI initially takes the user to the line of code matching the last FSA transition in this error trace.

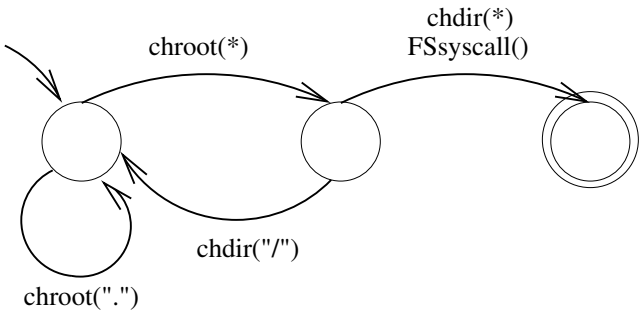

Figure 5. Our preliminary FSA to detect incorrect uses of `chroot()`. This FSA leads to a few false positives and is refined in Figure 6.

C Correct Use of `chroot()`

The `chroot()` system call is tricky to use properly, and historically some programs have been rendered vulnerable by a failure to follow safe usage patterns. To give some background, calling `chroot(dir)` is intended to confine the program within the portion of the filesystem under `dir`, the *jail*. This is implemented by the kernel as follows: after `chroot(dir)`, the program's accesses to `/` are mapped to `dir` in the underlying filesystem (so that `dir` acts as the root of the filesystem for this one process), and as a security check the program's accesses to `..` from within that directory are re-directed so that the program cannot escape from its jail. So long as the program starts within this jail, the `chroot()` call is irrevocable: even if the program is later penetrated through, say, a buffer overrun, the attacker will not be able to escape this jail or affect any part of the filesystem outside of the jail. Consequently, many security-sensitive programs use the `chroot()` system call to voluntarily give up unneeded privileges and thereby reduce the consequences from any unexpected security compromises.

However, there is a subtle pitfall here. Due to the way that most Unix kernels implement this functionality, if the program's current working directory is outside the jail at the time of the `chroot()` call, it is possible to escape the jail. Consider a program whose current directory is initially `/home/joe`, and that calls `chroot("/var/jail")`

but never changes its current directory. Suppose there is a vulnerability that allows an attacker to subsequently take control of the program, and the attacker executes `chdir("../..")`. This will succeed in changing the current directory to `/` (the real root of the filesystem, not the jail): directory traversal has never attempted to look up the name `..` from within the directory `/var/jail`, so the security checks are never triggered. At this point the attacker can modify the entire filesystem at will simply by using relative paths to refer to files outside the jail. In effect, this exploits a deficiency in the kernel's implementation of the `chroot()` system call.

One might think it would be easy to modify the kernel to eliminate this vulnerability, and indeed it would be, if not for one difficulty: this dangerous behavior has been enshrined in the POSIX standard as required for POSIX compliance. Consequently, most[8] Unix kernels retain this dangerous functionality.

This, then, is a nasty pitfall for the unwary programmer; if the programmer is not careful, her intent to confine an untrusted program can be foiled. To prevent this risk, the safe way to use `chroot()` is to make sure that the program's current working directory[9] is within the jail. The simplest pattern is to simply use

```
chroot(dir);
chdir("/");
```

in place of `chroot(dir)`.

Accordingly, we built a property to look for deviations from this safe pattern. Any program that calls `chroot()` and subsequently performs any filesystem-related call without first executing `chdir("/")` is declared in violation. We quickly discovered that we needed to modify the property to recognize error-handling code, like the following, as safe:

```
rv = chroot(dir);
if (rv < 0) {
  perror("chroot");
  exit(1);
}
chdir("/");
```

Therefore, as a special case, we modified the property so that any program trace that exits immediately after a `chroot()` is not considered an error, even if it fails to call `chdir("/")` before exiting. See Figure 5 for a graphical depiction of the corresponding FSA. We define *FSsyscall* to be any system call that interacts explicitly or implicitly with the file system. There are over 100, so we do not list them here.

[8]Not all Unices have allowed this dangerous behavior to remain. OpenBSD has decided to forsake standards compliance in favor of reducing the risk of attack.

[9]Similar care must be taken with any open file descriptors the program might retain across the `chroot()` call. Our property does not attempt to check this requirement.

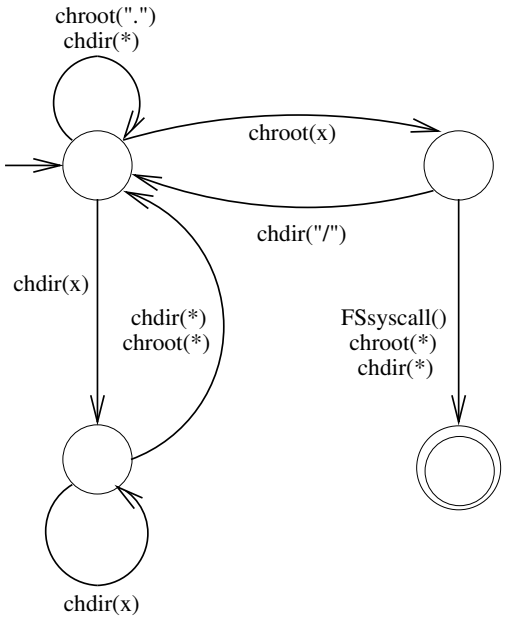

Figure 6. Our refined FSA for detecting `chroot()` vulnerabilities. This improves on the FSA in Figure 5 by reducing false positives.

After running this property across all Red Hat packages, we discovered several other common idioms for using `chroot()` that are safe. Because these idioms were not included in the aforementioned property, they led to 5 false positives. In particular, one safe idiom is

```
chdir(dir);
chroot(dir);
```

Similarly, another safe pattern is to move to the desired directory and then execute

```
chroot(".");
```

The obvious next step is to refine the FSA to allow all three of these patterns of usage. This refinement step turned out to be rather tricky, partially because `chdir(X); chroot(X)` is safe but `chdir(X); chroot(Y)` is not. We implemented this with a judicious use of pattern variables. See Figure 6.

With the aid of the refined property, we were able to confirm that none of the Red Hat packages contain unsafe usages of `chroot()`. There was only one false positive; see Figure 7. Arguably, the offending code segment should have been written more clearly, though the programmer has added a comment to clarify what is going on here. For comparison, there were 51 uses of `chroot()` spread across 30

```
if (guest || restricted) {
    if (chroot(pw->pw_dir) < 0)
        goto bad;
}
if (krb5_seteuid((uid_t)pw->pw_uid) < 0)
    goto bad;
if (guest) {
    /*
     * We MUST do a chdir() after the chroot. Otherwise
     * the old current directory will be accessible as "."
     * outside the new root!
     */
    if (chdir("/") < 0)
        goto bad;
} else {
    if (chdir(restricted ? "/" : pw->pw_dir) < 0) {
        if (chdir("/") < 0)
            goto bad;
    }
}
```

Figure 7. The only false positive for the `chroot()` property. This code, taken from Kerberos 5's ftpd program, is safe; however, MOPS is unable to prove its safety, since MOPS does not perform path-sensitive data flow analysis.

different packages in the Red Hat sources. The low false positive rate shows that our refined property could reasonably be applied as part of standard development and release processes.

D Preventing Format String Vulnerabilities

Some library functions are passed arguments containing format strings. For example, `printf()`'s first argument is a format string, and special sequences in it, such as `%d`, `%s`, etc., are replaced by values located on the stack. However, if the program passes an untrusted format string to `printf()`, a problem arises: the program will blindly attempt to retrieve values from the stack. This can lead to a compromise of security, as the user can specify format strings that cause `printf()` to write to arbitrary locations in memory. For example, suppose `buf` is a character buffer supplied by the user; then the call `printf(buf)` will introduce a security vulnerability. The correct usage is `printf("%s", buf)`. This type of vulnerability has been explored in previous work [17].

We detect potential format string vulnerabilities in a crude fashion: we identify all places in the code where format string functions may be supplied with non-constant character strings. Note that this introduces many false positives: for instance, `printf(buf)` might be safe if `buf` is a buffer known not to contain any format string sequences (for instance, it may have been initialized elsewhere to a value not under user control), but MOPS has no way of recognizing its safety in this case. Nonetheless, despite the presence of false positives, this simple property does represent one way to find all potential format string bugs in C code. We have built a FSA that identifies all format string functions and checks for non-constant arguments. Our FSA understands 61 different format string-related functions, and thereby generalizes and extends the set of functions checked for in Flawfinder, RATS, ITS4, and PScan.

After some experimentation, we soon encountered one important idiom with non-constant arguments to `printf` that is nonetheless recognizably safe. For instance, `printf(gettext(s))` is safe if `s` is a constant string, since we know the argument to `printf` will not contain a format string. Here we are assuming that the internationalization table is not maliciously constructed, hence the value `gettext(s)` can be trusted. We extended our FSA to recognize these additional safe programming idioms, thus reducing the number of false positives. With these improvements, we expect that our property will find at least as many bugs as previous tools, like RATS, ITS4, Flawfinder, and PScan, but with somewhat fewer false positives.

Nonetheless, when applying this property on all Red Hat 9 packages, the number of false positives still remains very

large. We have mainly used this example to illustrate the versatility of MOPS; our auditing is incomplete, and we have no plans to investigate all packages—partly due to the sheer magnitude of traces that need to be examined.

E Results

Tables 2, 3, 4, and 5 give further details on the results of our analysis of Red Hat 9 packages. For each package, we show the program affected, the number of warnings (counted in terms of the number of trace groups), and the number of real security bugs. For space reasons, our summary for the TOCTTOU property omits those packages with no real bugs.

Package	*Program*	*Warnings*	*Bugs*
binutils-2.13.90.0.18-9	ar	2	1
coreutils-4.5.3-19	chown	3	2
coreutils-4.5.3-19	chmod	2	1
coreutils-4.5.3-19	cp	2	1
dos2unix-3.1-15	dos2unix	4	2
ftpcopy-0.5.2-2	ftpcopy	8	3
gaim-0.59.8-0	gaim	3	3
joe-2.9.7-12	joe	1	1
jpilot-0.99.2-1	jpilot	2	1
initscripts-7.14-1	minilogd	1	1
inn-2.3.4-2	fastrm	1	1
isdn4k-utils-3.1-62	idsnlog	4	1
lrzsz-0.12.20-16	lsz	4	1
LPRng-3.8.19-3	checkpc	8	1
make-3.79.1-17	make	1	1
mc-4.6.0-4	mc	5	1
ncompress-4.2.4-33	compress	2	1
nmap-3.00-4	nmap	2	1
nmh-1.0.4-18	ali	1	1
pam-0.75-48	pam_console	5	1
postfix-1.1.11-11	postlock	2	1
postgresql-7.3.2-3	psql	4	1
rdist-6.1.5-26	rdist	18	3
rsh-0.17-14	rcp	4	2
strace-4.4.95-2	strace	1	1
symlinks-1.2-18	symlinks	4	2
talk-0.17-20	talkd	1	1
unix2dos-2.2-19	unix2dos	4	2
util-linux-2.11y-9	sln	2	1
zip-2.3-16	zip	7	1
Total		108	41

Table 2. Results: the TOCTTOU property.

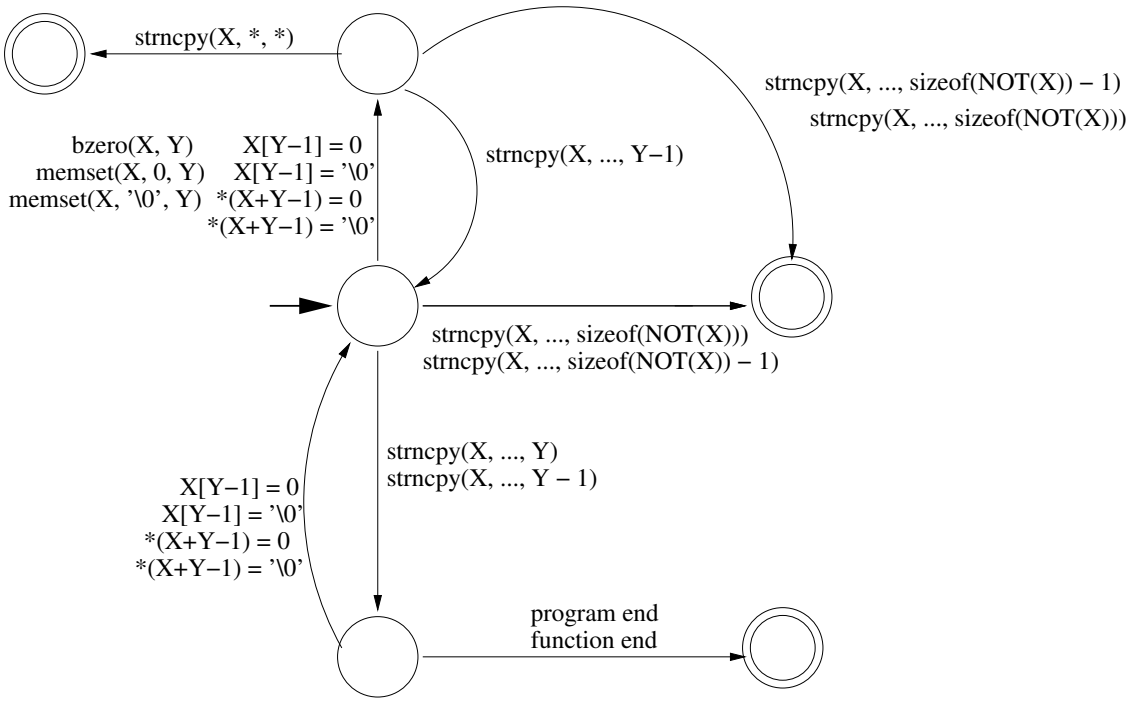

Figure 8. Our FSA for finding vulnerabilities arising from a misuse of *strncpy*.

Package	*Program*	*Warnings*	*Bugs*	*Exploits*
amanda-server-2.4.3-4	dumper	2	1	0
amanda-server-2.4.3-4	amcheck	4	2	0
at-3.1.8-33	at	2	1	0
gnuchess-5.02-11	gnuchess	6	1	0
httpd-2.0.40-21	suexec	2	1	0
inn-2.3.4-2	rnews	5	1	0
isdn4k-utils-3.1-62	userisdncntrl	1	1	0
kon2-0.3.9b-16	kon	3	1	0
lockdev-1.0.0-23	lockdev	1	1	0
mgetty-1.1.30-2	faxq-helper	5	4	0
ncpfs-2.2.1-1	ncplogin, ncpmap	9	1	0
rsh-0.17-14	rcp	2	1	0
slocate-2.6-8	slocate	3	1	0
sudo-1.6.6-3	sudo	3	2	0
util-linux-2.11y-9	chfn, chsh	3	1	0
uucp-1.06.1-47	cu	2	1	0
vixie-cron-3.0.1-74	crontab	3	1	0
Total		56	22	0

Table 3. Results: the standard file descriptor property.

Package	*Program*	*Warnings*	*Bugs*
anacron-2.3-25	anacron	1	0
binutils-2.13.90.0.18-9	ar	2	2
compat-gcc-7.3-2.96.118	texindex	1	0
cproto-4.6-15	cproto	1	0
inn-2.3.4-2	actsync	1	0
inn-2.3.4-2	ctlinnd	1	0
inn-2.3.4-2	innxmit	1	0
inn-2.3.4-2	rnews	1	0
inn-2.3.4-2	shrinkfile	1	0
imap-2001a-18	impad	1	1
isdn4k-utils-3.1-62	vbox	2	0
krb5-1.2.7-10	kprop	1	0
lha-1.14i-9	lha	1	0
libungif-4.1.0-15	gifinto	1	1
lrzsz-0.12.20-16	lrz	1	1
lv-4.49.4-9	lv	2	0
m4-1.4.1-13	m4	1	0
mailx-8.1.1-28	mail	1	0
mikmod-3.1.6-20	mikmod	1	1
mpage-2.5.3-3	mpage	1	0
ncurses4-5.0-11	tic	1	1
nmh-1.0.4-18	comp	2	0
patch-2.5.4-16	patch	1	0
patchutils-0.2.19-1	interdiff	1	0
pine-4.44-18	impad	1	0
pinfo-0.6.6-4	pinfo	1	0
psutils-1.17-19	psbook	1	0
rcs-5.7-20	co	2	0
rdist-6.1.5-26	rdistd	1	1
rdist-6.1.5-26	rdist	1	1
util-linux-2.11y-9	pg	2	0
vnc-3.3.3r2-47	imake	2	0
x3270-3.2.19-4	mkfb	1	0
Total		40	9

Table 4. Results: insecure creation of temporary files (excluding *mkstemp* template reuse; see Table 5).

Package	*Program*	*Warnings*	*Bugs*
byacc-1.9-25	yacc	5	3
chkconfig-1.3.8-1	ntsysv	1	0
dos2unix-3.1-15	dos2unix	4	2
dump-0.4b28-7	restore	1	0
file-3.39-9	file	1	1
gftp-2.0.14-2	gftp-gtk	1	1
ggv-1.99.97-2	ggv	1	1
gnome-libs-1.4.1..	gconfigger	2	1
gpm-1.19.3-27	gpm	1	0
isdn4k-utils-3.1-62	idsnrep	2	0
mars-nwe-0.99pl20-12	nwbind	2	0
mktemp-1.5-18	mktemp	1	1
mpage-2.5.3-3	mpage	1	1
nautilus-cd-burn..	nautilus-cd..	1	1
nkf-2.01-2	nks	1	1
ntp-4.1.2-0.rc1.2	ntp-genkeys	1	1
nvi-m17n-1.79-20..	nvi	2	2
openssh-3.5p1-6	sshd	1	0
pax-3.0-6	pax	2	2
quota-3.06-9	edquota	4	4
redhat-config-pr..	traditional	1	0
rwall-0.17-17	rwall	1	1
shadow-utils-4.0..	useradd	3	0
skkinput-2.06.3-3	skkinput	2	0
sndconfig-0.70-2	sndcondig	1	0
sylpheed-0.8.9-3	slypheed	1	0
xfig-3.2.3d-12	xfig	2	1
zebra-0.93b-1	bgpd	2	1
Total		48	25

Table 5. Results: re-use of the *mkstemp* template.