*RSE Interface:* This converts the CHECK instructions from the main processor into signals specific to the path-tracking module. This is done by tapping the Fetch\_out signal from the main pipeline. The *Fetch\_out* is one of the signals provided by the RSE framework [1]. Similarly, it converts signals from the path-tracking module into flags in the main processor. These are represented as special-purpose registers in the main processor

**Module Operation:** The operation of the path-tracking module for each of the primitive operations (executed in the main processor) is considered below:

CHECK instruction with emitEdge operation is executed in the main processor:

- RSE interface asserts the *emitEdge* signal and sends the basic block identifiers that constitute the edge in the *from* and *to* lines.
- The *from* and *to* identifiers are looked up in the *edge table* and the edge index corresponding to the edge is sent to the *state transition table*.
- The row corresponding to the edge is looked up in the *state transition table*.
- For each non-empty table-entry in the column corresponding to the checks, the states in the LHS of the transitions stored in the table entry are compared to the current state of the check in the *state vector*.
- If the states match, then the transition is fired and the state vector entry corresponding to the check is updated with the state in the RHS of the transition that matched.

CHECK instruction with the getState operation is executed in the main processor:

- RSE interface asserts the *getState* signal and sends the identifier of the check on the *checkID* line to the path-tracking module.
- The path tracking module looks up the state in the *state vector* and sends it to the RSE interface through the *currentState* line. This in turn is sent to the main processor and is returned as the value of the CHECK instruction (through a special register in the RSE).

*CHECK instruction with resetState operation is executed in the main processor***:** This is similar to the *getState* operation, but no value is returned to the RSE interface.

*Function calls/returns:* Since the technique tracks only intra-procedural paths, the *state vector* needs to be preserved across function calls and returns. This is done by pushing the *state vector* on a separate stack (different from the function call stack) along with the return address upon a function call and by popping the stack upon a return. The VRP generates code that uses special CHECK instructions to manipulate the stack on function calls/returns.

#### 4.8.1 Area Overheads

The area overheads for the hardware module are dominated by the three main components of the module presented in Section 0. The other components are mainly glue combinational logic and occupy negligible area. Table 21 presents the formulas used in estimating the size of the dominant hardware components. The size of each of these components depends on (1) the number of control-flow edges corresponding to state transitions (m), (2) the number of checks that must be tracked for the application (n) and, (3) the maximum number of transitions in each entry of the transition table because the table must be big enough to hold the biggest entry (k).

Hardware Component	Size (bits)	Explanation
Edge Table	m * 16 * 3	Each entry has 3 fields from, to and edgeIndex. Each fiels consists of 16 bits.
State Vector	n * 8	Each entry of the state vector consists of 8 bits, number of bits used for states
Transition Table	n * m * k * 16	Each state transition has two 8-bit fields to encode the starting and ending states

**Table 21: Formulas for estimating hardware overheads** 

Table 22 presents the values of m, n and k for each application as well as the number of bits occupied by each structure. The sizes of the hardware structures (in bits) are calculated based on Table 21.

App Name	m	n	k	Size (bits)
IntMM	10	21	3	1278
RealMM	10	21	3	1278
FFT	17	30	4	3096
Quicksort	19	29	5	3899
Bubblesort	5	11	1	383
Treesort	10	20	4	1440
Perm	16	27	1	1416
Queens	5	20	1	500
Towers	11	31	1	1117
Health	9	52	1	1316
Em3d	8	30	3	1344
Mst	17	33	10	6690
Barnes-Hut	43	118	6	33452
Tsp	9	48	2	1680
Average	14	37	4	4928

Table 22: Sizes of hardware structures (in bits)

The average number of bits stored by the hardware module is estimated to be 4928. This corresponds to less than 1 KB of storage space in the hardware. The application exhibiting the worst-case overhead (Barnes-hut) occupies 33452 bits, corresponding to less than 4KB of memory. This easily fits into a standard FPGA BRAM cell which has 5096KB of memory.

#### **4.8.2** Performance Overheads

The path-tracking module needs to be synchronized with the main processor only at the *getState* operation, and can execute asynchronously the rest of the time. Note that in our implementation of the path-tracking module the *getState* operation is simply a lookup in the state vector and takes constant time. Hence, *the getState operation takes constant time*. The *emitEdge, enterFunc* and *leaveFunc* operations can be buffered by the path-tracking module, while the application continues to execute on the main processor. These operations can then be performed asynchronously by the path-tracking module. We found that a buffer size of 1 sufficed to store the operations from the main processor.

#### 4.9 CONCLUSION

This chapter presented a technique to derive error detectors for protecting an application from data errors. The error detectors were derived automatically using compiler-based static analysis from the backward program slice of critical variables in the program. The slice is optimized aggressively based on specific control-paths in the application, to form a checking expression. At runtime, the control path executed by the program is tracked using specialized hardware and the corresponding checking expressions are executed.

# CHAPTER 5 FORMAL VERIFICATION OF ERROR DETECTORS

#### 5.1 INTRODUCTION

Error detection mechanisms are vital for building highly reliable systems. However, generic detection mechanisms such as exception handlers can take millions of processor cycles to detect errors in programs [14]. In the intervening time, the program can execute with the activated error and perform harmful actions such as writing incorrect state to the file-system. There has been significant work on efficiently placing [17, 40] and deriving [56, 109, 110] error detectors for programs. An important challenge is to enumerate the set of errors the mechanism fails to detect, either from a known set or an unknown set. Typically, verification techniques target the defined set of errors the detector is supposed to detect. However, one cannot predict the kinds of errors that may occur in the field, and hence it is important to evaluate detectors under arbitrary conditions.

Fault-injection is a well-established to evaluate the coverage of error detection mechanisms [19]. However, there is a compelling need to develop a formal framework to reason about the efficiency of error detectors as a complement to traditional fault injection. This chapter shows how this can uncover possible "corner cases" which may be missed by conventional fault injection due to its inherent statistical nature. While there have been formal frameworks, each addresses a specific error detection mechanism (for example replication [111]), and cannot be easily extended to general detection mechanisms.

*This chapter presents SymPLFIED, a framework for verifying error detectors in programs using symbolic execution and model-checking.* The goal of the framework is to expose error cases that would potentially escape detection and cause program failure. The focus is on transient hardware errors. The framework makes the following unique contributions:

- Introduces a formal model to represent programs expressed in a generic assembly language, and reasons about the effects of errors originating in hardware and propagating to the application without assuming specific detection mechanisms,
- Specifies the semantics of general error detectors using the same formalism, which allows verification of their detection capabilities,
- Represents errors using a single symbol, thereby coalescing multiple error values into a single symbolic value in the program. This includes both single- and multibit errors in the register file, main memory, cache, as well as errors in computation.

To the best of our knowledge, this is the first framework that models the effect of arbitrary hardware errors on software, independent of the underlying detection mechanism. It uses model checking [112] to exhaustively enumerate the consequences of the symbolic errors on the program. The analysis is completely automated and does not miss errors that might occur in a real execution. However as a result of symbolically abstracting erroneous values, it may discover errors that may not manifest in the real execution of the program i.e. false-positives. Previous work [113] has analyzed the effect of hardware errors on programs expressed in a high-level language (e.g. Java). Errors are modeled as bit flips in single data variable(s) in the program. While this is an important step, there are several limitations, namely (1) low-level hardware errors can affect multiple program variables as well as impact the program's control-flow, (2) errors in special-purpose registers such as the *stack pointer* are difficult to model in the high-level language, (3) Errors in the language runtime system (and libraries) cannot be modeled as they may be written in a different language. This chapter considers programs represented at the assembly language level. The value of using assembly language is that any low-level hardware error that impacts the program can be represented at the assembly language level (as shown in section 5.3.3). Further, the entire application, including runtime libraries is amenable to analysis at the assembly language level.

It can be argued that in order to really analyze the impact of hardware errors, we need to model systems at even lower levels, *e.g* the register-transfer level (RTL). However, the consequent state space explosion when analyzing the entire program at such low levels can impact the practicality of the model. *An assembly language representation is a judicious tradeoff between the size of the model and the representativeness of hardware errors that can be considered in the model.* 

In order to evaluate the framework, the effects of hardware transient errors are considered on a commercially deployed application, *tcas*. The framework identified errors that lead to a catastrophic outcome in the application, while a random fault injection experiment

did not find any catastrophic scenario in a comparable amount of time. The framework is also demonstrated on a larger program, *replace*, to find instances of incorrect program outcomes due to hardware transient errors.

#### 5.2 RELATED WORK

Prior literature related to this work is classified into the following categories:

**Error Detection:** Many error detection mechanisms have been proposed in the literature, along with formal proofs of their correctness [114, 115]. However, the verification methodology is usually tightly coupled with the mechanism under study. For example, Nicolescu et al. [116] proposes and verifies a control-flow checking technique by constructing a hypothetical program augmented with the technique and model-checks the program for missed detections. The program is carefully constructed to exercise all possible cases of the control-flow checking technique. It is non-trivial to construct such programs for other error-detection mechanisms.

Perry et al. [111] proposes the use of type-checking to verify the fault-tolerance provided by a specific error-detection mechanism namely, compiler-based instruction duplication. The paper proposes a detailed machine model for executing programs. The faults in the fault model (Single-Event Upsets) are represented as transitions in this machine model. The advantage of the technique is that it allows reasoning about the effect of low-level hardware faults on the whole program, rather than on individual instructions or data. However, the detection mechanism (duplication) is tightly coupled with the machine model, due to inherent assumptions that limit error propagation in the program and may not hold in non-duplicated programs.

Further, the type-checking technique in [111] either accepts or rejects a program based on whether the program has been duplicated correctly, but does not consider the consequences of the error on the program. As a result, the program may be rejected by the technique even though the error is benign and has no effect on the program's output.

**Symbolic execution** has been used for a wide variety of software testing and maintenance purposes [117]. The main idea in these techniques is to execute the program with symbolic values rather than concrete values and to abstract the program state as symbolic expressions. An example of a commercially deployed symbolic execution technique to find bugs in programs is Prefix [52]. However, Prefix assumes that the hardware does not experience errors during program execution.

A symbolic approach for injecting faults into programs was introduced by Larrson and Hahnle [113]. The goals of this approach are similar to ours, namely to verify properties of fault-tolerance mechanisms in the presence of hardware errors. The technique reasons on programs written in Java and considers the effect of bit-flips in program variables. However, a hardware error can have wide-ranging consequences on the program, including changing its control-flow and affecting the runtime support mechanisms for the language (such as the program stack and libraries). These errors are not considered by the technique.

Further, the technique presented in [113] uses theorem-proving to verify the errorresilience of programs. Theorem-proving has the intrinsic advantage that it is naturally symbolic and can reason about the non-determinism introduced by errors. However, as it stands today, theorem proving requires considerable programmer intervention and expertise, and cannot be completely automated for many important classes of programs.

**Program verification techniques** have been used to prove that a program's code satisfies a programmer-supplied specification [118]. The specification precisely outlines the expected result of the program given certain initial conditions. Typically, program verification techniques are geared towards finding software defects and assume that the hardware and the program environment are error-free. In other words, they prove that the program satisfies the specification *provided* the hardware platform on which the program is executed does not experience errors.

Further, program verification techniques operate on an abstract representation of the program (such as a state machine) extracted from the program code [72, 119]. The abstractions are derived based on the specific property being checked and cannot be used for evaluating the program under arbitrary hardware errors as such errors may not manifest in the abstraction.

Formal techniques have also been extensively applied to **microprocessor verification** [120]. The techniques attempt to prove that the implementation of the processor conforms to an architectural specification usually in the form of a processor reference manual.

Processor verification techniques focus on unmasking hardware design defects, as opposed to transient errors due to electrical disturbances or radiation.

**Soft-errors in Hardware:** Krautz et al. [121] and Seshia et al. [122] consider the effects of hardware transient errors (soft errors) on error-detection mechanisms implemented in hardware. While these techniques are useful for applications implemented as hardware circuits, it is not clear how the technique can be extended for reasoning about the effects of errors on programs. This is because programs are normally executed on general-purpose processors in which the manifestation of a low-level error is different from an error in an ASIC implementing the application.

**Summary:** The formal techniques considered in this section predominantly fall into the category of software-only techniques which do not consider hardware errors [118], or into the category of hardware-only techniques which do not consider the effects of errors on software [120]. Further, existing verification techniques are often coupled with the detection mechanism (e.g. duplication) being verified [111, 116].

Therefore, there exists no *generic* technique that allows reasoning about the effects of *arbitrary* hardware faults on software, and can be combined with an arbitrary fault model and detection technique(s). This is important for enumerating all hardware transient errors that would escape detection and cause programs to fail. Moreover, the technique must be *automated* in order to ensure wide adoption, and should not require programmer intervention.

This chapter attempts to answer the question: "Is it possible to develop a framework to reason about the effects of arbitrary hardware errors on applications in an automated fashion, in order to understand where error detection mechanisms fail in detecting errors?

## 5.3 APPROACH

This section, introduces the conceptual model of the SymPLFIED framework and also the technique used by SymPLFIED to symbolically propagate errors in the program. The fault-model used by the technique is also discussed.

#### 5.3.1 Framework

The SymPLFIED framework accepts a program protected with error detectors and enumerates all errors (in a particular class) that would not be detected by the detectors in the program. Figure 43 presents the conceptual design flow of the SymPLFIED framework.

**Inputs:** The inputs to the framework are (1) a program written in a target assembly language (e.g. MIPS), (2) error detectors embedded in the program code, and (3) a class of hardware errors to be considered (e.g. control-flow errors, register file errors).

**Assembly Language:** We define a generic assembly language in which programs are represented for formal analysis by the framework. Because the language defines a set of architectural abstractions found in many common RISC processor architectures, it is currentl portable across these architectures, with an architecture specific front-end. The assembly language has direct support for (1) Input/Output operations, so that programs can be analyzed independent of the Operating System (OS), and (2) Invocation of error detectors using special annotations, called CHECK, which allows detectors to be represented in line with the program.



Figure 43: Conceptual design flow of SymPLFIED

**Operation:** The program behavior is abstracted using a generic assembly language described in Section 5.5. This is automatically translated into a formal mathematical model that can be represented in the Maude system [34]. Since the abstraction is close to the actual program in assembly language it is sufficient for the user to formulate generic specifications, such as an incorrect program outcome or an exception being thrown. Such a low-level abstraction of the program is useful to reason about hardware errors. The formal model can then be rigorously analyzed under error conditions against the above specifications using techniques such as model-checking and theorem-proving. *In this* 

chapter, model-checking is used because it is completely automated and requires no programmer intervention.

**Outputs:** The framework uses the technique described in section 5.3.2 and outputs either of the following:

- 1. Proof that the program with the embedded detectors is resilient to the error class considered, OR
- 2. A comprehensive set of all errors belonging to the error class that evade detection and potentially lead to program failure (crash, hang or incorrect output).

Components: The framework consists of the following formal models,

- **Machine Model:** Models the formal semantics of the machine on which the program is to be executed (e.g. registers, memory, instructions etc.).
- Error Model: Specifies error classes and error manifestations in the machine on which the program is executed e.g. errors in the class *register errors* can manifest in any register in the machine.
- **Detector Model:** Specifies the format of error detectors and their execution semantics. It also includes the action taken upon detecting the error e.g. halting the program.

By representing all three models in the same formal framework, we can reason about the effects of errors (in the error model) on both programs, represented in the machine model and on detectors, represented in the detector model, in a unified fashion.

**Correctness:** In order for the results of the formal analysis to be trustworthy, the model must be provably correct. There are two aspects to correctness, namely,

- The model must satisfy certain desirable properties such as termination, coherence and sufficient completeness [34], AND
- 2. The model must be an accurate representation of the system being modeled.

The first requirement can be satisfied by formally analyzing the specification using automated checking tools for each desirable property listed above. This is obtained almost for free by expressing the model using Maude as formal checking tools are available to check the conformance of the model to the properties [123].

However, the second requirement is much harder to ensure as it cannot be checked by formal tools and is usually left to the model creator. We have attempted to validate the model by rigorously analyzing the behavior of errors in the model and comparing them with the behavior of the real system (Section **5.6.3**).

#### **5.3.2** Symbolic Fault Propagation

The SymPLFIED approach represents the state of all erroneous values in the program using the abstract symbol *err*. The *err* symbol is propagated to different locations in the program during execution using simple error propagation rules (shown in section 5.5.2). The symbol also introduces non-determinism in the program when used in the context of comparison and branch instructions or as a pointer operand in memory operations. Because the same symbol is used to represent all erroneous values in the program, the approach distinguishes program states based on where errors occur rather than on the nature of the individual error(s). As a result, it avoids state explosion and can keep track of all possible places in the program the error may propagate to starting from its origin. However, because errors in data values are not distinguished from each other, the set of error states corresponding to a fault is over-approximated. This can result in the technique finding erroneous program outcomes that may not occur in a real execution. For example, if an error propagates from a program variable A to another variable B, the variable B's value is constrained by the value of the variable A. In other words, given a concrete value of A after it has been affected by the error, the value of B can be uniquely determined due to the error propagating from A to B.

The SymPLFIED technique on the other hand, would assign a symbolic value of *err* to both variables, and would not capture the constraint on *B* due to the variable *A*. As a result, it would not be able to determine that the value in register *B* even when given the value in register *A*. This may result in the technique discovering spurious program outcomes. Such spurious outcomes are termed *false-positives*.

While SymPLFIED may uncover false-positives, it will never miss an outcome that may occur in the program due to the error (in a real execution). This is because SymPLFIED systematically explores the space of all possible manifestations of the error on the program. Hence, the technique is *sound*, meaning it finds all error manifestations, but is not always *accurate*, meaning that it may find false-positives.

Soundness is more important than accuracy from the point of view of designing detection mechanisms, as we can always augment the set of error detectors to conservatively protect against a few false-positives (due to the inaccuracies introduced).

While a small number of false-positives can be tolerated, it must be ensured that the technique does not find too many false-positives as the cost of developing detectors to protect against the false-positives can overwhelm the benefits provided by detection. The SymPLFIED technique uses a custom constraint solver to remove false-positives in the search-space. The constraint solver also considerably limits state space explosion and quickly prunes infeasible paths. More details may be found in Section 5.5.2.

#### 5.3.3 Fault Model

The fault-model considered by SymPLFIED includes transient errors in memory/registers and computation.

- Errors in memory/registers are modeled by replacing the contents of the memory location or register by the symbol *err*. *No distinction is made between single- and multi-bit errors*.
- Errors in computation are modeled based on where they occur in the processor pipelin*e and* how they affect the architectural state as shown in Table 23.
- Errors in processor control-logic (such as in the register renaming unit) are not considered by the fault-model.

The reason it is possible to represent such a broad class of errors in the model is because the program is represented in assembly language, which makes the elements of its state explicit to the analysis framework.

Fault origin	Error Symptom	Conditions under which Modeled	Modeling procedure	
Instruction Decoder	One of the fields of an instruction is corrupted	One valid instruction is converted to another valid instruction	Instructions writing to a destination (e.g., add) - change the output target Instructions with no target (e.g., nop) – replace with instructions with targets (e.g. add) Instructions with a single destination (e.g.add) – replace with instruction with	err in the original and new targets (register or memory)         err in the new wrong target (register or memory)         err in the original target location (register or memory)
Address or Data Bus	Data read from memory, cache or register file is corrupted	Single and multiple bit errors in the bus during instruction execution	no target (e.g. <i>nop</i> ) Errors in register data bus	<i>err</i> in source register(s) of the current instruction <i>err</i> in target registers of <i>load</i>
			Error in memory bus	instructions to the location err in target register of <i>load</i> instructions to the location
Processor Functional Unit	Functional unit output is corrupted	Single and multiple bit errors in registers/me mory	Functional Unit output to register or memory	<i>err</i> in register or memory file being written to by the current instruction
Instruction Fetch Mechanism	Errors in the fetch unit	Single or multiple bit errors in PC or instruction	Fetch from an erroneous location due to error in <i>PC</i>	<i>PC</i> is changed to an arbitrary but valid code location
			Error in instruction while fetching	Modeled as Decode Errors

Table 23: Computation error categories and how they are modeled by SymPLFIED

#### 5.3.4 Scalability

As in most model-checking approaches, the exhaustive search performed by SymPLFIED can be exponential in the number of instructions executed by the program in the worst case. In spite of this limitation, model-checking techniques have been successfully scaled to large code-bases such as operating system kernels and web-servers [72, 119]. These approaches consider only parts of the system that are relevant to the property being verified. The relevant code portions are typically extracted by static analysis. However, static analysis is not very useful for dealing with runtime errors that may occur in hardware.

However, the error detection mechanisms in the program can be used to optimize the state space exploration process. For example, if a certain code component protected with detectors is proved to be resilient to all errors of a particular class, then such errors can be ignored when considering the space of errors that can occur in the system as a whole. This lends itself to a hierarchical or compositional approach, where first the detection mechanisms deployed in small components are proved to protect that component from errors of a particular class, and then inter-component interactions are considered. This is an area of future investigation.

#### 5.4 EXAMPLES

This section illustrates the SymPLFIED approach in the context of an application that calculates the factorial of a number shown in Figure 2. The program is represented in the generic assembly language presented in Section 5.3.1.

#### 5.4.1 Error Injection

We illustrate our approach with an example of an injected error in the program shown in Figure 43. Assume that a fault occurs in register \$3 (which holds the value of the loop counter variable) in line 8 of the program after the loop counter is decremented (*subi* \$3 \$3 1). The effect of the fault is to replace the contents of the register \$3 with *err*. The loop back-edge is then executed and the loop condition is evaluated by (*setgt* \$5 \$3 \$4). Since \$3 has the value *err* in it, it cannot be determined if the loop condition evaluates to true or false. Therefore, the execution is forked so that the loop condition evaluates to make the set of true in one case and to false in the other case. The *true* case exits immediately and prints

the value stored in \$2. Since the error can occur in any loop iteration, the value printed can be any of the following: *1!*, *2!*, *3!*, *4!*, *5!*. All these outcomes are found by SymPLFIED.

1	ori \$2 \$0 #1	initial product p = 1	
2	read \$1	read i from input	
3	mov \$3, \$1	-	
4	ori \$4 \$0 #1	for comparison purposes	
loop:	setgt \$5 \$3 \$4	start of loop	
6	beq \$5 0 exit	loop condition : \$3 > \$4	
7	mult \$2 \$2 \$3	$p = p * i$	
8	subi \$3 \$3 #1	i = i - 1	
9	beq \$0 #0 loop	loop backedge	
exit:	prints "Factorial =	= "	
11	print \$2		
12	halt		

Figure 44: Program to compute factorial in MIPS-like assembly language

The *false* case continues executing the loop and the *err* value is propagated from register \$3 to register \$2 due to the multiplication operation (*mul* \$2 \$2 \$3). The program then executes the loop back-edge and evaluates the branch condition. Again, the condition cannot be resolved as register \$3 is still *err*. The execution is forked again and the process is repeated ad-infinitum. In practical terms, the loop is terminated after a certain number of instructions and the value *err* is printed, or the program times out<sup>18</sup> and is stopped.

**Complexity:** Note that in order for a physical fault-injection approach to discover the same set of outcomes for the program as SymPLFIED, it would need to inject into all possible values (in the integer range) into the loop counter variable. This can correspond to  $2^k$  cases in the worst case, where *k* is the number of bits used to represent an integer. In

<sup>&</sup>lt;sup>18</sup> We assume that a watchdog mechanism is present in the program to monitor for infinite loops and hangs.

contrast, SymPLFIED considers at-most (n+1) possible cases, in this example, where *n* is the number of iterations of the loop. This is because each fork of the execution at the loop condition results in the *true* case exiting the loop and the program. In the general case though, SymPLFIED may need to consider  $2^n$  possible cases. However, by upperbounding the number of instructions executed in the program, the growth in the searchspace can be controlled.

**False-positives:** In the example, not all errors in the loop counter variables will cause the loop to terminate early. For example, an error in the higher-order bits of the loop counter variable in register \$3 may still cause the loop condition (\$3 > \$4) to be *false*. However, SymPLFIED would conservatively assume that both the *true* and *false* cases are possible, as it does not distinguish between errors in different bit-positions of variables. Note that in practice, false-positives were not a major concern, as shown in section **5.6.2**.

### 5.4.2 Error Detection

We now discuss how SymPLFIED supports error-detection mechanisms in the program. Figure 45 shows the same program in Figure 44, augmented with error detectors. Recall that detectors are invoked through special CHECK annotations as explained in Section 5.3.1. The error detectors together with their supporting instructions (*mov* instruction in line 8) are shown in bold.

```
ori $2 $0 #1
                                  --- initial product p = 1
2
        read $1
                                  --- read i from input
3
        mov $3, $1
4
        ori $4 $0 #1
                                  --- for comparison purposes
       setgt $5 $3 $4
                                  --- start of loop
loop:
6
         beq $5 0 exit
7
        check ($4 < $3)
8
         mov $6, $2
9
        mult $2 $2 $3
                             ---- p = p * i
10
          check ($2 >= $6 * $1)
11
         subi $3 $3 #1
                             ---- i = i - 1
         beq $0 #0 loop
12
                            --- loop backedge
       prints "Factorial = "
exit:
          print $2
14
15
          halt
```

Figure 45: Factorial program with error detectors inserted

The same error is injected as before in register \$3 (the new line number is 11). As shown in Section 5.4.1, the loop back-edge is executed and the execution is forked at the loop condition (\$3 > \$4).

The *true* case exits immediately, while the *false* case continues executing the loop. The *false* case "remembers" that the loop condition (\$3 < \$4) is false by adding this as a constraint to the search. The *false* case then encounters the first detector that checks if (\$4 < \$3). The check always evaluates to *true* because of the constraint and hence does not detect the error.

The program continues execution and the error propagates to \$2 in the *mul* instruction. However, the value of \$2 from the previous iteration does not have an error in it, and this value is copied to register \$6 by the *mov* instruction in line 8. Therefore, when the second detector is encountered within the loop (line 10), the LHS of the check evaluates to *err* and the RHS evaluates to (\$6 \* \$1), which is an integer.

The execution is forked once again at the second detector into *true* and *false* cases. The *true* case continues execution and propagates the error in the program as before. The *false* case of the check throws an exception and the detector fails, thereby detecting the error.

The constraints for the *false* case, namely,  $(\$6 * \$3 \ge \$6 * \$1)$  are also remembered. Based on this constraint, as well as the earlier constraint  $(\$3 \ge \$4)$ , the constraint-solver deduces that the second detector will detect the error if and only if the fault in register \$3causes it to have a value greater than the initial value read from the input (stored in register \$1).

The programmer can then formulate a detector to handle the case when the error causes the value of register \$3 to be lesser than the original value in register \$1. Therefore, the errors that evade detection are made explicit to the programmer (or to an automated mechanism) who can make an informed decision about handling the errors.

The error considered above is only one of many possible errors that may occur in the program. These errors are too numerous for manual inspection and analysis as done in this example. Moreover, not all these errors evade detection in the program and lead to program failure.

The main advantage of SymPLFIED is that it can quickly isolate the errors that would evade detection and cause program failure from the set of all possible transient errors that can occur in the program. It can also show an execution trace of how the error evaded detection and led to the failure. This is important in order to understand the weaknesses in existing detection mechanisms and improve them.

#### 5.5 IMPLEMENTATION

We have implemented the SymPLFIED framework using the Maude rewriting logic system.

**Rewriting logic** is a general-purpose logical framework for specification of programming languages and systems [124].

**Maude** is a high-performance reflective language and system supporting both equational and rewriting logic specification and programming for a wide range of applications [34]. *The main advantage of Maude is that it allows a wide variety of formal analysis techniques to be applied on the same specification.* 

**Supporting Tools:** In order to make programs for existing architectures compatible with SymPLFIED, we provide a facility (through means of a *Perl* script) to translate programs written directly in the target architecture's assembly language into SymPLFIED's assembly language. In theory, any architecture can be supported but for now we support only the MIPS instruction set. We also built a query generator to explore the behavior of the program under common hardware error categories. Note that while the SymPLFIED framework can support arbitrary error classes, pre-defined error categories allow programmers to verify the resilience of their programs without having to write complex specifications (or any specifications).

In this section, we describe the details of the machine, detector and error models and show how the resilience of programs to hardware errors can be verified through exhaustive search i.e. bounded model-checking.

#### 5.5.1 Machine Model

This section describes the machine model for executing assembly language programs using Maude.

**Equations and Rules:** As far as possible, we have used equations instead of rewrite rules for specifying the models. The main advantage of using equations is that Maude performs rewriting using equations much faster than using rewrite rules. However, equations must be deterministic and cannot accommodate ambiguity. The machine model is completely deterministic because for a given instruction sequence, the final state can be uniquely determined in the absence of errors. Therefore the machine model can be represented entirely using equations. However, the error model is non-deterministic and hence requires rewrite rules.

**Assumptions:** The following assumptions are made by the machine model when executing a program.

- An attempt to fetch an instruction from an invalid code address results in an "illegal instruction" exception being thrown. The set of valid addresses is defined at program load time by the loader.
- 2. Memory locations are defined when they are first written to (by store instructions). An attempt to read from undefined memory location results in an "illegal address" exception being thrown. It is assumed that the program loader initializes all locations prior to their first use in the program.
- 3. Program instructions are assumed to be immutable and hence cannot be overwritten during execution.
- 4. Arithmetic operations are supported only on integers and not on floating point numbers.

**Machine State**: The central abstraction used in the machine model is the notion of *machine state*, which consists of the mutable components of the processor's structures. The machine state is carried from instruction to instruction in program execution order, with each instruction optionally looking up and/or updating the state's contents. The machine state is obtained by concatenating one or more of the machine elements in a single 'soup' of entities. For example, the soup, PC(pc) regs(R) mem(M) input(In) *output(out)* represents a machine state in which the (1) current program counter is denoted by *pc*, (2) register file is denoted by *R*, (3) memory is denoted by *M* and (4) input and output streams are *in* and *out* respectively.

**Execute Sub-Model:** We consider example instructions from each instruction class and illustrate the equations used to model them. These equations are defined in the *execute* sub-model and use primitives defined in other sub-models (e.g. the *fetch* primitive).

1. Arithmetic Instruction: Consider the execution of the addi instruction, which adds the value19 v to the register given by rs and stores the results in register rd. In the equation given below, the <\_,\_> operator represents the machine state obtained by executing an instruction (given by the first argument) on a machine state (given by the second argument). C represents the code of the program and is written outside the state to enable faster rewriting by Maude (as it is assumed to be immutable). The {\_,\_} groups together the code and the machine-state. The elements of the machine

<sup>&</sup>lt;sup>19</sup> The term *value* is used to refer to both integers and the *err* symbol

state in the equations are composable, and hence can be matched with a generic symbol *S* representing the "rest of the state". This allows new machine-state elements can be added without modifying existing equations.

 $eq \{ C, < addi rd rs v, PC(pc) regs(R) S > \} = \{ C, < fetch(C, pc), PC(next(pc)) regs(R[rd] < -R[rs] + v) S > \}.$ 

2. Branch Instructions: Consider the example of the *beq rs*, *v*, *l* instruction, which branches to the code label *l* if and only if the register *rs* contains the constant value *v*. The equation for *beq* is similar to the equation for the *addi* operation except that it uses the in-built if-then-else operator of Maude. Note the use of the *isEqual* primitive rather than a direct == to compare the values of the register *rs* and the constant value *v*. This is because the register *rs* may contain the symbolic constant *err* and hence needs to be resolved accordingly (by the error model).

 $eq \{ C, < beq rs \lor l, pc(PC) regs(R) S > = if isEqual(R[rs], \lor) then \{ C, < fetch(C, pc), PC(next(pc)) regs(R) S > \} else \{ C, < fetch(C, l), PC(l) regs(R) S > \} fi.$ 

- **3.** Load/Store Instructions: Consider the example of the *ldi rt, rs, a* which loads the value in the memory location at the address given by adding the offset *a* to the value in the register *rs*. However, the load address needs to be checked for validity before loading the value. This is done by the *isValid* primitive (defined in the *Memory Submodel*).
  - $eq \{ C, < ldi rt rs a, PC(pc) regs(R) mem(M) S > = if (isValid(R[rd] + a, M)) then \{ C, < fetch(C, pc), C(next(pc)) mem(M) regs(R[rt] <- M[a + R[rs]]) > \} else \{ C, < throw "Illegal addr", PC(next(pc)) mem(M) regs(R) > \} fi.$

4. Input/Output Operations: Input and output operations are supported natively on the machine since the operating system is not modeled. An example is the print instruction whose equation is as follows:

 $eq \{ C, < print rs, PC(pc) regs(R) output(O) S > \} = \{ C, < fetch(C, pc), PC(next(pc)) regs(R) output(O << R[rd]) S > \}.$ 

5. Special Instructions: These instructions are responsible for starting and stopping the program. e.g. *halt* and *throw* instructions to terminate the program. The halt instruction transforms the super-state prior to their execution into a machine state in order to facilitate the search for final solutions by the model-checker (section 5.5.4). Its equation is given by:

$$eq \{ C, < halt, PC(pc) S > \} = PC(done) S.$$

#### 5.5.2 Error Model

The overall approach to error injection and propagation was discussed in Section 5.3.2, but in this section we discuss the implementation of the approach using rewriting logic in Maude. The implementation of the error model is divided into five sub-models as follows:

**Error Injection Sub-Model:** The error-injection sub-model is responsible for introducing symbolic errors into the program during its execution. The injector can be used to inject the *err* symbol into registers, memory locations or the program counter when the program reaches a specific location in the code. This is implemented by adding a breakpoint mechanism to the machine model described in Section 5.5.1 The choice of

which register or memory location to inject into is made non-deterministically by the injection sub-model using rewrite rules.

**Error Propagation Sub-Model:** Once an error has been injected, it is allowed to propagate through the equations for executing the program in the machine model. The rules for error propagation are also described by equations as shown below. In the equations that follow, *I* represents an integer.

 $eq \ err + err = err . eq \ err + I = err . eq \ I + err = err .$   $eq \ err - err = err . eq \ err - I = err . eq \ I - err = err .$   $eq \ err * I = if \ (I==0) \ then \ 0 \ else \ err \ fi .$   $eq \ err / I = if \ (I==0) \ then \ throw \ "div-zero" \ else \ err \ fi .$   $eq \ err * err = if \ isEqual(err, 0) \ then \ throw \ "div-zero" \ else \ err \ fi .$   $eq \ err / err = if \ isEqual(err, 0) \ then \ throw \ "div-zero" \ else \ err \ fi .$ 

In other words, any arithmetic operation involving the *err* value also evaluates to *err* (unless it is multiplied by 0). Note also how the divide-by-zero case is handled.

**Comparison Handling Sub-Model:** The rules for comparisons involving one or more *err* values are expressed as rewrite-rules as they are non-deterministic in nature. For example, the rewrite rules for the *isEqual* operator used in section 5.5.1 are as follows:

rl isEqual(I, err) => true . rl isEqual(I, err) => false .
rl isEqual(err, err) => true . rl isEqual(err, err) => false .

The comparison operators involving err operands evaluate to either true or false nondeterministically. This is equivalent to forking the program's execution into the true and false cases. However, once the execution has been forked, the outcome of the comparison is deterministic and subsequent comparisons involving the same unmodified locations must return the same outcome (otherwise false-positives will result). This can be accomplished by updating the state (after forking the execution) with the results of the comparison. In the *true* case of the *isEqual* primitive, the location being compared can be updated with the value it is being compared to. However, the *false* case is not as simple, as it needs to "remember" that the location involved in the comparison is not equal to the value it is being compared with. The same issue arises in the case of non-equality comparisons, such as *isGreaterThan*, *isLesserThan*, *isNotGreaterThan* and *isNotLesserThan*.

The *constraint tracking and solving* sub-model remembers these constraints and determines if a set of constraints is satisfiable, and if not, truncates the state-space exploration for the case corresponding to the constraint. This helps avoid reporting false-positives.

**Constraint Tracking and Solving Sub-Model:** A new structure called the *ConstraintMap* is added to the machine state in Section 5.5.1. The *ConstraintMap* structure maps each register or memory location containing *err* to a set of constraints that are satisfied by the value in the location. An example of a set of constraints for a location is the following: *notGreaterThan(5) notEqualTo(2) greaterThan(0)*. This indicates that the location can take any integer value between 0 and 5 excluding 0 and 2 but including

5. The constraints for a location are updated whenever a comparison is made based on the location if and only if it contains the value *err*. Constraints are also updated by arithmetic and logic operations in the program.

For a given location, it may not be possible to find a value that satisfies all its constraints simultaneously. Such constraints are deemed un-satisfiable and the model-checker can terminate the search when it comes to a state with an un-satisfiable set of constraints (such a state represents a false-positive). The constraint solver determines whether a set of constraints is un-satisfiable and eliminates redundancies in the constraint-set.

**Memory- and Control Handling Sub-Model:** Memory and Control errors are also handled non-deterministically using rewrite rules as follows:

*Errors in jump or branch targets:* The program either jumps to an arbitrary (but valid) code location or throws an "illegal instruction" exception.

*Errors in pointer values of loads:* The program either retrieves the contents of an arbitrary memory location or throws an "illegal-address" exception.

*Errors in pointer values of stores:* The program either overwrites the contents of an arbitrary memory location, or creates a new value in memory.

#### 5.5.3 Detector Model

Error detectors are defined as executable checks in the program that test whether a given memory location or register satisfies an arithmetic or logical expression. For example, a detector can check if the value of register (5) equals the sum of the values in the register (3) and memory location (1000) at a given program counter location. If the values do not match, an exception is thrown and the program is halted.

In our implementation, each detector is assigned a unique identifier and the CHECK instructions encode the identifier of the detector they want to invoke in their operand fields. The detectors themselves are written outside the program, and the same detector can be invoked at multiple places within the program's code.

We assume that the execution of a detector does not fail i.e. the detectors themselves are free of errors.

A detector is written in the following format:

det (ID, Register Name or Memory Location to Check, Comparison Operation, Arithmetic Expression )

- 1. The arguments of the detector are as follows:
- 2. The first argument of the detector is its identifier.
- 3. The second argument is the register or memory location checked by the detector.
- The third argument is the comparison operation, which can be any of ==, =/=, >,
   <, <= or >=.
- 5. The final argument is the arithmetic expression that is used to check the detector's register or memory location and is expressed in the following format:

Expr ::= Expr + Expr | Expr - Expr | Expr \* Expr | Expr / Expr | (c) | (Reg Name) | \* (memory address)Using the above notation, the detector introduced earlier would be written as:

$$det(4, \ \$(5), ==, (\ \$3) + \ \ast(1000)).$$

The equations for the detector's execution are independent of the equations in the machine model, and hence are not affected by errors introduced in the machine other than those that are present in the registers or memory locations used in the detector's expression. Execution of a detector also updates the constraints for the location being checked in the *ConstraintMap* structure described in section 5.5.2.

#### 5.5.4 Model-checking

The exhaustive *search* feature of Maude is used to model-check programs [123]. The aim of the search command is to expose interesting "outcomes" of the program caused by errors in a particular category. The "outcome" is a user-defined function on the machine state described in Section 5.5.1 and must be specified in the *search* command. For example, the following search command obtains the set of executions of the program that will print a value of *err* under all single errors in registers (one per execution).

#### $search \ regErrors(\ start(program, first, \ detectors)) = >! \ (S:MachineState) \ such \ that \ (\ output(S) \ contains \ err).$

The *search* command systematically explores the search space in a breadth-first manner starting from the initial state and obtaining all final states that satisfy the user-defined predicate, which can be any formula in first-order logic. The programmer can query how specific final states were obtained or print out the search graph, which will contain the entire set of states that have been explored by the model checking. This can help the programmer understand how the injected error(s) lead to the outcome(s) printed by the search.

**Termination:** In the absence of errors, most programs can be modeled as finite-space systems provided (1) they terminate after a finite amount of time or (2) they perform repetitive actions without terminating but revisit states. However, errors can cause the state space to become infinitely large, as the program may loop infinitely due to the error, never revisiting earlier states. In practice, this is impossible, since the program data is physically represented as bits and there are only a finite number of bits available in a machine. However, the state space would be so large that it is practically impossible to explore in full.

In order to ensure that the model-checking terminates, the number of instructions that is allowed to be executed by the program must be bounded. This bound is referred to as the *timeout* and must be conservatively chosen to encompass the number of instructions executed by the program during all possible correct executions (in the absence of errors). After the specified number of instructions is exceeded, a "timed out" exception is thrown and the program is halted. We assume that the processor has a watchdog mechanism.

#### 5.6 CASE STUDY

We have implemented SymPLFIED using Maude version 2.1. Our implementation consists of about 2000 lines of uncommented Maude code split into 35 modules. It has 54 rewrite rules and 384 equations.

This section reports our experience in using SymPLFIED on the *tcas* application [125], which is widely used as an advisory tool in air traffic control for ensuring minimum vertical separation between two aircrafts and hence avoid collisions. The application

consists of about 140 lines of C code, which is compiled to 913 lines of MIPS assembly code, which in turn is translated to 800 lines of SymPLFIED's assembly code (by our custom translator). In the later part of this section, we describe how we apply SymPLFIED on the *replace* program of the Siemens program suite [51] to understand the effects of scaling to larger programs.

*tcas* takes as input a set of 12 parameters indicating the positions of the two aircrafts and prints a single number as its output. The output can be one of the following values: 0, 1 or 2, where 0 indicates that the condition is unresolved, 1 indicates an upward advisory and 2 indicates a downward advisory. Based on these advisories, the aircraft operator can choose to increase or decrease the aircraft's altitude.

#### 5.6.1 Experiment Setup

Our goal is to find whether a transient error in the register file during the execution of *tcas* can lead to the program producing an incorrect output (in this case, an advisory). We chose an input for *tcas* in which the *upward advisory* (value of 1) would be produced under error-free execution.

We directed SymPLFIED to search for runs in which the program did not throw an exception and produced a value other than 1 under the assumption of a single register error in each execution. The search command is identical to the one shown in section 5.5.4.

This constitutes about (800 \* 32) possible injections, since there are 32 registers in the machine, and each instruction in the program is chosen as a breakpoint. In order to reduce

the search space, at each breakpoint, only the register(s) used by the instruction was injected. This ensures that the fault is activated in the program.

In order to ensure quick turn-around time for the injections, they were started on a cluster of 150 dual-processor AMD Opteron machines. The search command is split into multiple smaller searches, each of which sweeps a particular section of the program code looking for errors that satisfy the search conditions. The smaller searches can performed independently by each node in the cluster, and the results pooled together to find the overall set of errors. The maximum number of errors found by each search task was capped at 10, and a maximum time of 30 minutes was allotted for task completion (after which the task was killed).

In order to validate the results from SymPLFIED, we augmented the Simplescalar simulator [50] with the capability to inject errors into the source and destination registers of all instructions, one at a time. For each register we injected three extreme values in the integer range as well as three random values, so that a representative sample of the errors in each value can be considered.

#### 5.6.2 SymPLFIED Results

For the *tcas* application, we found only one case where an output of 1 is converted to an output of 2 by the fault injections. This can potentially be catastrophic as it is hard to distinguish from the correct outcome of *tcas*. None of the other injections found any other such case. We also found cases where (1) *tcas* printed an output of 0 (unresolved) in place of 1, (2) the output was outside the range of the allowed values printed by *tcas* and
(3) numerous cases where the program crashed. We do not report these cases as *tcas* is only an advisory tool and the operator can ignore the advisory if he or she determines that the output produced by *tcas* is incorrect.

We also found violations in which the value is computed correctly but printed incorrectly. We do not consider these cases as the output method may be different in the commercial implementation of *tcas*.

**Running Time:** Of the 150 search tasks started on the cluster, 85 tasks completed within the allotted time of 30 minutes. The remaining 65 tasks did not complete in the allotted time (as the timeout chosen was too large). We report results only from the tasks that completed. Of the 85 tasks that completed, 70 tasks did not find any errors that satisfy the conditions in the search command (as either the error was benign or the program crashed due to the error). These 70 tasks completed within 1 minute overall.

The time taken by the 15 completed tasks that found errors satisfying the search condition, (including the catastrophic outcome) is less than 4 minutes, and the average time for task completion is 64 seconds. Even without considering the incomplete tasks we were able to find the catastrophic outcome for *tcas*, shown below.

Initially, we were surprised by the unusually low number of catastrophic failures reported in *tcas*. However, closer inspection revealed that the code has been extremely wellengineered to prevent precisely these kinds of error from resulting in catastrophic failures. The *tcas* application has been extensively verified and checked for safety violations by multiple studies [126-128]. Nevertheless, the fact that SymPLFIED found this failure at all is testimony to its comprehensive evaluation capabilities. Further, this failure was not exposed by the injections performed using Simplescalar. In order to understand better the error that lead to *tcas* printing the incorrect value of 2, we show an excerpt from the *tcas* code in Figure 46.

```
int alt_sep_test()
  enabled = High_Confidence && (Own_Tracked_Alt_Rate <= OLEV) && (Cur_Vertical_Sep > MAXALTDIFF);
  tcas_equipped = Other_Capability == TCAS_TA;
  intent_not_known = Two_of_Three_Reports_Valid &&
                                                          (Other_RAC == NO_INTENT);
  alt_sep = UNRESOLVED;
  if (enabled && ((tcas_equipped && intent_not_known) || !tcas_equipped)) {
         need_upward_RA = Non_Crossing_Biased_Climb() && Own_Below_Threat();
         need_downward_RA = Non_Crossing_Biased_Descend() && Own_Above_Threat();
         if (need_upward_RA && need_downward_RA)
           alt_sep = UNRESOLVED;
         else if (need_upward_RA)
           alt_sep = UPWARD_RA;
         else if (need_downward_RA)
           alt_sep = DOWNWARD_RA;
         else
           alt_sep = UNRESOLVED;
    return alt_sep;
```

Figure 46: Portion of *tcas* code corresponding to error

**Optimizations:** In order to reduce the number of states explored by the model-checker, we inject errors only into the registers used in each instruction of the program. Further, we inject the error just before the instruction that uses the register, in order to ensure fault activation. The effect of the injection is equivalent to injecting the register at an arbitrary code location so that the error is activated at the instruction.

The code shown in Figure 46 corresponds to the function *alt\_sep\_test*, which tests the minimum vertical separation between two aircrafts and returns an advisory. This function in turn calls the function *Non\_Crossing\_Biased\_Climb()* and the *Own\_Above\_Threat()* function to decide if an upward advisory is needed for the aircraft. It then checks if a downward advisory is needed by calling the function *Non\_Crossing\_Biased\_Descend()* 

and the function *Own\_Below\_Threat()*. If neither or both advisories are needed, it returns the value *0* (unresolved). Otherwise, it returns the advisory computed in the function.

The error under consideration occurs in the body of the called function  $Non\_Crossing\_Biased\_Climb()$  and corrupts the value of register \$31which holds the function return address. Therefore, instead of control being transferred to the instruction following the call to the function  $Non\_Crossing\_Biased\_Climb()$  in  $alt\_sep\_test()$ , the control gets transferred to the statement  $alt\_sep = DOWNWARD\_RA$  in the function. This causes the function to return the value 2 instead of the value 1, which is printed by the program. We have verified that the error exposed above corresponds to a real error and is not a false-positive by injecting these faults into the augmented Simplescalar simulator.

Note that the above error occurs in the stack, which is part of the runtime support added by the compiler. Hence, in order to discover this error, we need a technique like SymPFLIED that can reason at the assembly language (or lower) level.

#### 5.6.3 SimpleScalar Results

We performed over 6000 fault-injection runs on the *tcas* application using the modified Simplescalar simulator to see if we can find the catastrophic outcome outlined above. We ensured that both SymPLFIED and Simple-scalar were run for the same time to find these outcomes. The SymPLFIED injections were run with 150 tasks, and each completed task took a maximum time of 4 minutes. This constitutes 10 hours in total. We were able to perform 6000 automated fault-injection experiments with Simplescalar in that time. The results are summarized in column 2 of Table 24.

Program Outcome	Percentage		
	# faults = 6253	# faults = 41082	
0	1.86% (117)	2.33% (960)	
1	53.7% (3364)	56.33% (23143)	
2	0% (0)	0% (0)	
Other	0.5% (29)	1.0% (404)	
Crash	43.4% (2718)	40.43% (16208)	
Hang	0.4% (25)	0.8% (327)	

Table 24: SimpleScalar fault-injection results

Table 24 shows that even though we injected exhaustively into registers of all instructions in the program, Simplescalar was unable to uncover even a single scenario with the catastrophic outcome of '2', whereas the symbolic error injection performed by SymPLFIED was able to uncover these scenarios with relative ease. This is because in order to find an error scenario using random fault injections, not only must the error be injected at the right place in the program (for example, register *\$31* in the *Non\_Crossing\_Biased\_Climb* function), but also the right value must be chosen during the injection (for example, the address of the assignment statement must be chosen in the *alt\_sep\_test* function in Figure 46. Otherwise the program may crash due to the error or the error may be benign in the program.

We also extended the SimpleScalar based fault injection campaign to inject 41000 register faults to check if such an injection discovers errors causing the catatrophic outcome. The injection campaign completed in 35 hours but was still unable to find such an error. The results of this extended set of injections in shown in column 3 of Table 24.

# 5.6.4 Application to Larger Programs

In order to evaluate the effectiveness of the formal analysis as we scale to larger applications, we analyzed the *replace* program using SymPLFIED. *replace* is the largest of the Siemens benchmarks[51], used extensively in software testing. The *replace* program matches a given string pattern in the input string and replaces it with another given string. The code translates to about 1550 lines of assembly code spanning 22 functions. Table 25 lists some key functions.

 Table 25: Important functions in replace

makepat	Constructs pattern to be matched from input reg exp
getccl	Called by makepat when scanning a '[' character
dodash	Called by getccl for any character ranges in pattern
amatch	Returns the position where pattern matched
locate	Called by amatch to find whether the pattern appears at a string index

Using the same experimental setup as described in Section 5.6.1, we ran SymPLFIED on the *replace* program to find all single register errors (one per execution) that lead to an incorrect outcome of the program. The overall search was decomposed into 312 search tasks.

**Results:** Of these 202 completed execution within the allotted time of 30 minutes. In 148 of the completed search tasks, either the error was benign or the program crashed due to the error, while 54 of the search tasks found error(s) leading to incorrect outcome. We consider the execution trace of an example error.

**Example Scenario:** An input parameter to the dodash function that holds the delimiter (']') for a character range was injected. An erroneous pattern is constructed, which leads

to a failure in the pattern match. As a result, the program returns the original string without the substitution. The analysis completed in an average of 4 minutes where no erroneous solutions where found. For the injection runs that found an erroneous outcome the analysis took an average of 10 minutes.

# 5.7 CONCLUSION

This chapter presented SymPLFIED a modular, flexible framework for performing symbolic fault-injection and evaluating error-detectors in programs. We have implemented the SymPLFIED framework for a MIPS-like processor using the Maude rewriting logic engine. We demonstrate the SymPLFIED framework on a widely-deployed application *tcas*, and use it to find a non-trivial case of a hardware transient error that can lead to catastrophic consequences for the *tcas* system. We also demonstrate the technique on the *replace* program to illustrate its scalability.

# CHAPTER 6 FORMAL VERIFICATION OF ATTACK DETECTORS

### 6.1 INTRODUCTION

Insider threats have gained prominence as an emerging and important class of security threats [129-131]. An insider is a person who is part of the organization and either steals secrets or subverts the working of the organization by exploiting hidden system flaws for malicious purposes. This chapter considers *application-level* insider attacks. For example, an insider may load a malicious plugin into a web browser that overwrites the address bar with the address of a phishing website. Or a disgruntled programmer may plant a logical flaw in a banking application that allows an external user to fraudulently withdraw money. Both are examples of how a trusted insider can compromise an application and subvert it for malicious purposes.

We define an application-level insider attack as one in which a malicious insider attempts to overwrite one or more data items in the application, in order to achieve a specific attack goal. The overwriting may be carried out by exploiting existing vulnerabilities in the application (e.g. buffer overflows), by introducing logical flaws in the application code or through malicious third-party libraries. It is also possible (though not required) to launch insider attacks from a malicious operating system or higher-privileged process. Application-level insider attacks are particularly insidious because, (1) by attacking the application an insider can evade detection by mimicking its normal behavior (from the point of view of the system), and (2) to attack the application, it is enough for the insider to have the same privilege as that of the application, whereas attacking the network or operating system may require super-user privileges.

Before defending against insider attacks, we need a model for reasoning about insiders. Previous work has modeled insider attacks at the network and operating system (OS) levels using higher-level formalisms such as attack graphs [132-134] and process calculi [135]. However, modeling application-level insider attacks requires analysis of the application's code as an insider has access to the application and can hence launch attacks on the application's implementation. Higher-level models are too coarse grained to enable reasoning about attacks that can be launched at the application code level. Further, higher-level models typically require application vulnerabilities (if present) to be identified up-front in order to reason about insider attacks on the system.

This chapter introduces a technique to formally model application-level insider attacks on the application code expressed in assembly language. The advantage of modeling at the assembly code-level is that the assembly code includes the program, its libraries, and any state added by the compiler (e.g. stack pointer, return addresses). Therefore, all *softwarebased* insider attacks on the application can be modeled at the assembly-code level.

The proposed technique uses a combination of symbolic execution and model checking to systematically enumerate *all* possible insider attacks in a given application corresponding to an attack goal. The technique can be automatically deployed on the application's code and no formal specifications need to be provided other than generic specifications about the attacker's end goal(s) (with regard to the application's state or final output).

The value of the analysis performed by the proposed technique is that it can expose nonintuitive cases of insider attacks that may be missed by manual code inspection. This is because the technique exhaustively considers corruptions of data items used in the application (under a given input), and enumerates all corruptions that lead to a successful attack (based on the specified attack goal). Thus, it is able to identify *all vulnerable data items* in the application corresponding to the attack goal. The results of the analysis can be used to guide the development of defense mechanisms (eg. assertions) to protect the application.

We have implemented the proposed technique as a tool, *SymPLAID*, which directly analyzes MIPS-based assembly code. The tool identifies for each attack, (1) The program point at which the attack must be launched, (2) The data item that must be overwritten by the attacker, and (3) The value that must be used for overwriting the data item in order to carry out the attack.

SymPLAID is built as an enhancement of our earlier tool, SymPLFIED [136], used to evaluate the effect of transient errors on the application. SymPLFIED also builds a formal model of the application at the assembly code level. However, SymPLFIED groups individual errors into a single abstract class (*err*), and considers the effect of the entire class of errors on the program. This is because in the case of randomly occurring errors, we are more interested in the propagation of the error rather than the precise set of circumstances that caused the error.

195

In contrast, security attacks are launched by an intelligent adversary and hence it is important to know precisely what values are corrupted by the attacker (and how the corruption is carried out) in order to design efficient defense mechanisms against the attack(s). Therefore, SymPLAID was built from the ground up to emphasize precision in terms of identifying the specific conditions for an attack. Thus, rather than abstracting the attacker's behavior into a single class, the effect of each value corruption is considered individually, and its propagation tracked in the program. While this may appear to sacrifice scalability, the gains due to the extra precision in terms of evaluating fewer program forks outweigh the losses (see Section 6.4.4 for details).

The chapter makes the following key contributions:

- Introduces a formal model for reasoning about application-level insider attacks at the assembly-code level,
- Shows how application-level insiders may be able to subvert the execution of the application for malicious purposes,
- Describes a technique to automatically discover *all* possible insider attacks in an application using symbolic execution and model checking on the application code,
- Demonstrates the proposed techniques using a case-study drawn from the OpenSSH application[137], and finds all possible insider attacks, including several non-intuitive attacks that may be missed by simple, manual inspection.

# 6.2 INSIDER ATTACK MODEL

This section describes the attack model for insider attacks and an example scenario for an insider attack. The example scenario is considered in more detail in Section 6.3.

### 6.2.1 Characterization of Insider

**Capabilities**: The insider is a part of the application and has unfettered access to the program's address space. This includes the ability to both read and write the program's memory and registers. However, we assume that the insider cannot modify the program's code, which is reasonable since in most programs the code segment is marked read-only. An attacker may get into the application and become an insider in the following ways:

- By a logical loophole in the application planted by a disgruntled or malicious programmer,
- Through a malicious (or buggy) third-party library loaded into the address space of the application,
- By exploiting known security loopholes such as buffer overflow attacks and planting the attack code,
- By overwriting the process's registers or memory from another process (with higher privilege) or debugger,
- Through a security vulnerability in the operating system or virtual machine (if present)

In each of the above scenarios, the insider can corrupt the values of either memory locations or registers while the application is executing. The first three scenarios only require the insider to have the same privileges as the applications, while the last two require higher privileges.

**Goal**: The attacker's goal is to subvert the application to perform malicious functions on behalf of the attacker. However, the attacker wants to elude detection or culpability (as far as possible), so the attacker's code may not directly carry out the attack, but may instead overwrite elements of the program's data or control in order to achieve the attacker's aims. From an external perspective, it will appear as though the attack originated due to an application malfunction, and hence the attack code will not be blamed. Full execution replay may be able to find the attack [138], but it incurs considerable time and resource overheads. Therefore, the attacker can execute code to overwrite crucial elements of the program's data or control elements.

It is assumed that the attacker does not want to crash the application, but wants to subvert its execution for some malicious purpose. The attack is typically launched only under a specific set of inputs to the program (known to the attacker), and the input sequence that launches the attack is indistinguishable from a legitimate input for the program. Even if the insider is unable to launch the attack by himself/herself, he/she may have a colluding user who supplies the required inputs to launch the attack. Note that the colluding user does not need to have the same privileges as the insider in order to launch the attack.

198

### 6.2.2 Attack Scenario

Figure 47 shows an example attack scenario where the insider has planted a "logic bomb" in the application which is triggered under a specific set of inputs. The bomb could have been planted by the insider through the first, second or third scenario considered in section 6.2.1. .Normal users are unlikely to accidentally supply the trigger sequence and will be able to use the application without any problems. However, a colluding user knows about the bomb and supplies the trigger sequence as input. Perimeter based protection techniques such as firewalls will not notice anything amiss as the trigger sequence is indistinguishable from a regular input for all practical purposes. However, the input will trigger the bomb in the application thereby launching the security attack on behalf of the insider.



Figure 47: Attack scenario of an insider attack

# 6.2.3 **Problem Definition**

The problem of attack generation from the insider's point of view may be summed up as follows: "If the input sequence to trigger the attack is known (AND) the attacker's code

is executed at specific points in the program, what data items in the program should be corrupted and in what way to achieve the attack goal?"

This chapter develops a technique to automatically discover conditions for insider attacks in an application given (i) the inputs to trigger the attack (e.g. a specific user-name as input), (ii) the attacker's objective stated in terms of the final state of the application (e.g. to allow a particular user to log in with the wrong password) and (iii) the attacker's capabilities in terms of the points from which the attack can be launched (e.g. within a specific function). The analysis identifies both the target data to be corrupted and what value it should be replaced with to achieve the attacker's goal.

To facilitate the analysis, the following assumptions are made about the attacker by the technique.

- Only one value can be corrupted, but the corrupted value can be any valid value. This assumption ensures that the footprint of the attack is kept small and is hence easier to evade detection (from a defense technique)
- 2. The corruption is only allowed at fixed program points. This assumption reflects the fact that an insider may be able launch their attacks only at fixed program points.

We are working on relaxing these assumptions to consider attackers with higher capabilities or privileges. However, as we show in Section 6.3.1, even under these assumptions, a malicious insider can mount a significant number of attacks.

200

# 6.3 EXAMPLE CODE AND ATTACKS

This section considers an example code fragment to illustrate the attack scenario in Section 6.2.2. This is motivated by the OpenSSH program [137]. We consider the real OpenSSH application in Section 6.5. The example is also used in Section 6.5 to demonstrate the operation of the SymPLAID tool.

Figure 48 shows an example code fragment containing the *authenticate* function. The authenticate function copies the value of the system password into the *src* buffer and the value entered by the user into the *dest* buffer (in both cases it reads them into the *tmp* buffer first to validate the values). It then compares the values in the *src* and *dest* buffers and if they match, it returns the value 1 (authenticated). Otherwise it returns the value 0 (unauthenticated) to the calling function.



Figure 48: Code of authenticate function

# 6.3.1 Insider Attacks

We first take the attacker's perspective in coming up with insider attacks on the code in Figure 48. The attacker's goal is to allow a colluding user<sup>20</sup> to be validated even if he/she has entered the wrong password. The following assumptions are made in this example, for simplicity of explanation:

<sup>&</sup>lt;sup>20</sup> The colluding user may be the same person as the attacker (who wants to evade detection), but we distinguish between these two roles in this chapter.

- The attack can be invoked only within the body of the *authenticate* function.
- The attacker can overwrite the value of any register or local variable, but not global variables and heap buffers (due to practical limitations such as not knowing the exact address of globals and dynamic memory).
- The attack points are at function calls within the *authenticate* function, i.e., the arguments to any of the functions called by the *authenticate* function may be overwritten prior to the function call.

Table 26 shows the set of all possible attacks the attacker could launch in the above function. The first column shows the program point at which the attack is launched, the second column shows the variable to overwrite and the third column shows the value that should be written to the variable. The fourth column explains the attack in more detail. A particularly interesting attack found is presented in row 6 of Table 26, where the *dest* argument of the *strncpy* function was set to overlay the *src* string in memory. This replaces the first character of the *src* string with '\0', effectively converting it to a NULL string. The *dest* string also becomes NULL as the *dest* buffer is not filled by the *strncpy* function. As a result, the two strings will match when compared and the *authenticate* function will return '1' (authenticated).

As Table 1 shows, discovering all possible insider attacks manually (by inspection) is cumbersome and non-trivial even for the modestly sized piece of code that is considered in Figure 48. Therefore, we have developed a tool to generate the attacks automatically -

SymPLAID. Although the tool works on assembly language programs, we have shown the program as C-language code in Figure 48 for simplicity.

Program Point	Variable to be corrupted	Corrupted value of variable	Comments/Explanation	
	dest	src buf	The src buffer is compared with itself	
	src	dest buf	The dest buffer is compared with itself	
strncmp point (line 5)	src	tmp buf	The dest buffer is compared with the tmp buffer which contains the same string	
	len	<= 0	The strncmp function terminates early and returns 0 (the strings are identical)	
atmony point	temp	src buf	This copies the string in the source buffer to the destination buffer, thereby ensuring that the strings match	
(line 4)	dest	srcBuf – strlen(buf)	This writes a '\0' character in the src buffer, effectively converting it to a empty string. The dst buffer is also empty as it is not initialized, and hence the strings match.	
	temp	dest buf	The value in the two buffer is left unchanged and is	
readInput point (line 3)	temp	Any unused location in memory	copied to the dst buffer. This value is also stored in the src buffer and hence the strings match.	

 Table 26: Insider attacks on the authenticate function

We have validated the attacks shown in Table 26 using the GNU debugger (gdb) to corrupt the values of chosen variables in the application on an AMD machine running the Linux operating system. All the attacks shown in Table 26 were found to be successful i.e. they led to the user being authenticated in spite of providing the wrong password.

#### 6.3.2 Defense Techniques

We now take the defender's perspective in designing protection mechanisms for insider attacks.

The attacks in Table 26 consist of both "obvious attacks" as well as surprising corner cases. It can be argued that finding obvious attacks is not very useful as they are likely to be revealed by manual inspection of the code. However, the power of the proposed technique is that it can reveal *all* such attacks on the code, whereas a human operator may

miss one or more attacks. This is especially important from the developer's perspective, as *all* the security holes in the application need to be plugged before it can be claimed that the application is secure (as all the attacker needs to exploit is a single vulnerability). Moreover, the ability to discover corner-case attacks is the real benefit of using an automated approach.

Below we discuss some examples of detection mechanisms for the example presented in Figure 48. The mechanisms are designed based on the attacks discovered in Table 26.

- We insert a check before the call to the *strncmp* function. In particular, we check that the *src* and *dest* buffers of the *strncmp* function do not overlap with each other or with the *temp* buffer. We also check whether the length argument is greater than 0. This prevents attacks in rows 1 to 4 of Table 26. Note that the check is stronger than necessary.
- We insert a check after the call to the *readInput* function in line 3 to ensure that the *temp* buffer is non-empty. This prevents attacks in the rows 7 to 8 of Table 26.
- We insert a check before the call to the *strncpy* function to ensure that neither the *temp* buffer nor the *dest* buffer overlap with the *src* buffer.

If any of the above conditions is violated, the application is aborted and an attack is detected. The insider cannot corrupt the values in both the checks and the program as only one value in the application is allowed to be corrupted (as per our assumptions). Figure 49 shows the code in Figure 48 with the above checks inserted. The checks are represented as *assert* statements.

```
int authenticate(void* src, void* dest, void* temp, int len){
    1: readInput(temp);
    2: strncpy(src, temp, len)
    3: readInput(temp);
    assert( isNotEmpty(temp) );
    assert( noOverlap(temp, src) and noOverlap(temp, dest) )
    4: strncpy(dest, temp, len);
    assert( noOverlap(src, dest) and noOverlap(src, temp) );
    assert( len > 0 );
    5: if (! strncmp(dest, src, len) ) return 1;
    return 0;
}
```

Figure 49: Code of authenticate function with assertions

# 6.4 TECHNIQUE AND TOOL

As mentioned in the previous section, enumerating insider attacks by hand is cumbersome and non-trivial. Therefore, automating the discovery of insider attacks is essential. This section describes the key techniques used in the automation and the design of a tool to perform the discovery.

# 6.4.1 Symbolic Execution Technique

We represent an insider attack as a corruption of data values at specific points in the program's execution i.e. attack points. The attack points are chosen by the program developer based on knowledge of where an insider can attack the application. For example, all the places where the application calls an untrusted third-party library are attack points as an insider can launch an attack from these points. In the worst-case, every instruction in the application can be an attack point.

The program is executed with a known (concrete) input, and when one of the specified execution points is reached, a single variable<sup>21</sup> is chosen from the set of all variables in

<sup>&</sup>lt;sup>21</sup> We use the term variable to refer to both registers and memory locations in the program. This includes stack, heap and static data.

the program and assigned to a symbolic value (i.e. not a concrete value). The program execution is continued with the symbolic value. The above procedure is repeated exhaustively for each data value in the program at each of the specified attack points. This allows enumeration of all insider attacks on a given program.

The key technique used to comprehensively enumerate insider attacks is *symbolic execution-based model checking*. This means that the program is executed with a combination of concrete values and symbolic values, and model-checking is used to "fillin" the symbolic values as and when needed. Symbolic values are treated similar to concrete values in arithmetic and logical computations performed in the system. The main difference is in how branches and memory accesses based on expressions involving symbolic values are handled as follows:

When a branch decision involving a symbolic expression is reached, the program is forked – one fork executes the branch assuming that the branch decision is true, and the other fork executes the branch assuming that the branch decision is false. The branch decision is added as a constraint to the program state, and the symbolic expression is evaluated based on the constraint. In case the solution converges to a single value, all symbolic expressions in the state are replaced with their concrete values. Otherwise, the constraint is added to a pool of constraints and the program's execution is continued. The global pool of constraints is maintained in parallel with the concrete program state, and updated on each branch executed by the program. When a memory read (write) involving a symbolic expression in the address operand is encountered, the entire memory space of the system is scanned<sup>22</sup> and each location in memory is considered to be a potential target for the memory read (write). For each potential target, the execution of the program is forked and the symbolic expression is assigned to be equal to one of the addresses in memory. The value read (written) is assigned to the value stored in the corresponding address. As a result, the symbolic expression will evaluate to a unique value, which is then substituted in all symbolic expressions in the state. Thus, a memory access with a symbolic expression as the pointer operand converts the state into one in which all values are concrete.

Symbolic expressions are also used to represent indirect control transfers in the program (through a function pointer, for example). Indirect control transfers are treated similar to memory accesses through symbolic pointer expressions. In other words, each code location is treated as a potential target for the indirect branch, and the execution is forked with the constraint for the expression added to the fork.

For each program fork encountered above, the model checker checks whether (1) The fork is a viable one, based on the past constraints of the symbolic expressions, and (2) whether the fork leads to a desired outcome (of the attacker). If these two conditions are satisfied, the model checker will print the state of the program corresponding to the fork i.e. attack state.

<sup>&</sup>lt;sup>22</sup> The exhaustive search may incur significant overheads. Section 6.6.5 presents ways to make this search less performance intensive.

### 6.4.2 SymPLAID Tool

The symbolic execution technique described in the previous section has been implemented in an automated tool – SymPLAID (*Symbolic Program Level Attack Injection and Detection*). This is based on our earlier tool, SymPLFIED, used to study the effect of transient errors on programs [136]. The differences between SymPLAID and SymPLFIED are explained in Section 6.4.3.

SymPLAID accepts the following inputs: (1) an assembly language program along with libraries (if any), (2) a set of pre-defined inputs for the program, (3) a specification of the desired goal of the attacker (expressed as a formula in first-order logic) and (4) a set of attack points in the application. It generates a comprehensive set of insider attacks that lead to the goal state.

For each attack, SymPLAID generates both the location (memory or register) to be corrupted as well as the value that must be written to the location by the attacker. Figure 50 shows the conceptual view of SymPLAID from a user's perspective.



Figure 50: Conceptual view of SymPLAID's usage model

SymPLAID directly parses and interprets assembly language programs for a MIPS processor. The current implementation supports the entire range of MIPS instructions, including (1) arithmetic/logical instructions, (2) memory accesses (both aligned and unaligned) and (3) branches (both direct and indirect). However, it does not support system calls. The lack of system call support is compensated for by the provision of native support for input/output operations. Floating point operations are also not considered by SymPLAID. This is reasonable as floating-point operations are not typically used by security-critical code in the majority of applications.

SymPLAID is implemented using Maude, a high-performance language and system that supports specification and programming in rewriting logic [34]. SymPLAID models the execution semantics of an assembly language program using both equations and rewriting rules. Equations are used to model the concrete semantics of the machine, while rewriting rules are used for introducing non-determinism due to symbolic evaluation.

#### 6.4.3 Differences with SymPLFIED

This section discusses the differences between the analysis performed by the SymPLFIED and SymPLAID tools. The first column of Table 27 shows an example code fragment (in a MIPS-like assembly language). The state of each register in the program as determined by SymPLFIED and SymPLAID (after executing the instruction in the row) is shown in the second and third columns of Table 27 respectively.

Assume that the value in register \$2 has been corrupted (either by a transient error or by an insider attack) in instruction 2. Both SymPLAID and SymPLFIED represent the value

in register \$2 using the abstract symbol *err*. Until this point, the state of the program in both SymPLFIED and SymPLAID is the same. Then instruction 3 is executed, which subtracts 5 from the value in register \$2. At this point, the states diverge: SymPLFIED represents the value in register \$5 also with the symbol *err*, thereby approximating the dependencies among the value. On the other hand, SymPLAID represents the value in register \$5 by the symbolic expression (*err* – 5), which is more precise.

Code	SymPLFIED	SymPLAID
[ 1   movi \$3, #(10) ]	\$3 = 10	
[2   addi \$2, 0, #(err)]	2 = err	
[3   subi \$5, \$2, #(5)]	\$5 = err	\$5 = err - 5
[ 4   muli \$4, \$(5), #(2) ]	4 = 2 * err	4 = 2 * (err - 5)
[5   bne \$4, \$3, 7 ]	4 == 1  or  0?	4 = 1  or  0?
[ 6   print \$5 ]	2 = err	\$2 = 10
	3 = 10,	\$3 = 10
	\$4 = 10	\$4 = 10
	\$5 = err	\$5 = 5
	output: err	output: 5

Table 27: Example code illustrate SymPLFIED and SymPLAID

The execution then continues on to instruction 4, which multiplies the value in register \$5 with a constant 2. SymPLFIED once again approximates the value in register \$5 with the symbol *err*, whereas SymPLAID stores the symbolic expression 2 \* (err - 5) in register \$5.

Finally, execution reaches instruction 5, which checks if the value in register \$2 is not equal to the constant 10. If so, the program branches to location 7 and bypasses instruction 6. Otherwise, it executes instruction 6 which prints the value stored in register \$5.

In the case of both SymPLFIED and SymPLAID, the execution is forked at instruction 5. This is because the value in register \$4 contains a symbolic expression involving *err* in the case of SymPLAID, and the symbol *err* in the case of SymPLFIED. In both cases, the value cannot be uniquely evaluated. Hence, the tool must consider both the case where the equality holds and the case where it does not by forking the program's execution. Let us consider the fork where the equality holds and the branch is not taken. In this case, the control reaches statement 6 and the value in register \$6 is printed by the program. SymPLFIED and SymPLAID would print a different value at this instruction.

In the case of SymPLFIED, when the branch in instruction 5 is not taken, the value in register \$4 is set to 10. No other changes are made to the registers. Hence, when instruction 6 is reached, register \$5 contains the value *err*, which is printed.

In the case of SymPLAID, when the branch is not taken, the value of register \$4 is set to 10 as done in the case of SymPLFIED. However, the assignment to register \$4 triggers a wave of updates in the system. This is because the symbolic expression in register \$4, which is (err - 5) \* 2 is set equal to 10. The value of *err* is then uniquely determined to be 5 by solving the above equation. The values of registers \$2, \$4 and \$5 are updated based on the solved value to be 10, 10 and 5 respectively. This effectively converts all symbolic expressions in the state to concrete ones and is an example of condition 1 in section 6.4.1. Hence, when instruction 6 is reached, register \$5 contains the value *5*, which is printed.

The above example illustrates how SymPLAID is able to achieve higher precision than SymPLFIED by tracking dependencies among the corrupted values in registers and memory. *As a result, it is able to isolate the value(s) that must be injected by an attacker* 

211

*to achieve their attack goal.* In this example, if the attacker wants to make the program print the value 5, he/she must overwrite the value in register \$2 at instruction 2 with the value 5. The attack is automatically discovered by SymPLAID but not by SymPLFIED. The other differences with SymPLFIED are that SymPLAID supports a wider range of instructions (e.g. unaligned loads and stores) and has a more precise constraint solver.

#### 6.4.4 Precision and Scalability

SymPLAID maintains precise dependencies both in terms of arithmetic and logical constraints and solves them using a custom constraint solver. However, calls to the constraint solver can be expensive, and are minimized as follows:

- SymPLAID keeps track of symbolic expressions as part of the application state, and does not solve them until a decision point is reached, viz., branches, memory accesses and indirect control transfers.
- SymPLAID maintains simple linear constraints in a separate constraint map, and can identify infeasible states (such as the same erroneous location being assigned to multiple values) without invoking the constraint solver.
- Finally, SymPLAID replaces symbolic states with concrete states at the earliest opportunities, thereby reducing calls to the constraint solver and also minimizing the number of forks in the program.

There are two cases where SymPLAID performs approximations to conserve space. These are as follows: First, SymPLAID can only solve linear constraints. In practice, few security critical branches or memory access operations involve non-linear constraints and hence this does not result in false-positives. The second approximation occurs when an unaligned memory access is performed in SymPLAID. In order to conserve space, SymPLAID stores values in memory as integral values over the entire word, and hence cannot model corruptions of individual bytes in a word. This can result in loss of precision leading to false-positives.

# 6.5 DETAILED ANALYSIS

This section illustrates how SymPLAID identifies conditions for successful insider attacks in the context of the example considered in Section 6.3. The code shown in Figure 51 is a modified version of the MIPS assembly code for the example in Figure 49.

In Figure 51, instructions and labels are expressed within '[', ']', and a | character separates the instruction from its label. Comments are prefixed with a --- and can follow the instruction. For ease of analysis, we generated simplified versions of the standard library functions, but with the same functionality as the original functions.

Figure 51 shows the assembly code of the *strncmp* function and an excerpt from the *authenticate* function that calls the *strncmp* function. In this section, we will consider how SymPLAID analyzes the effect of data value corruptions introduced in the *authenticate* function to generate the set of attacks in the first three rows of Table 26.

Consider the attacks that can be launched at the call site of the *strncmp* function in instruction 44 of the *authenticate* function. This is the *attack point* in this example. We

assume that the insider can only corrupt values in registers. The following discussion considers only the cases that lead to successful attacks.

Strncmp: String comparison routine
[0   ldo \$(16) #(4) \$(esp) ] load the src
[1   ldo \$(17) #(8) \$(esp) ] load the dest
[2   ldo \$(18) #(12) \$(esp) ] load the length
[3   movi \$(1) #(1) ] result to return
[4   bltzi \$(18) #(15) ] while length > 0
[5   lbu \$(19) #(0) \$(16) ] load of src
[6   lbu \$(20) #(0) \$(17) ] load of dest
[7   seq \$(21) \$(19) \$(20) ] is *src==*dest ?
[8   beqii \$(21) #(0) #(15) ] if not equal, break
[9   beqii \$(19) #(0) #(16) ]if end of string, break
[ 10   addi \$(16) \$(16) #(1) ] increment src
[ 11   addi \$(17) \$(17) #(1) ] increment dest
[12   subi \$(18) \$(18) #(1) ] decrement length
[ 13   beqii \$(0) #(0) #(4) ] loop backedge
[ 14   movi \$(1) #(0) ] store in register 1
[ 15   return ] return
Authenticate Function Excerpt:
[21   ldo \$(1) #(tmpAddr) \$(0) ] Retrieve temp
[22   ldo \$(2) #(srcAddr) \$(0) ] Retrieve the src
[23   ldo  (3) #(lengthAddr) (0) ] Retrieve the length
$\begin{bmatrix} 24 \mid \text{sto } \$(1) \#(4) \$(\text{esp}) \end{bmatrix}$ Push parameters on stack
[25   sto \$(2) #(8) \$(esp) ]
[26   sto \$(3) #(12) \$(esp) ]
[27   call #(strncpyLoc)] call the string copy function
$\begin{bmatrix} 28 \mid \text{Ido } \$(1) \#(\text{tmpAddr}) \$(0) \end{bmatrix}$ load the temp buf
$[29   \text{sto } (1) \# (4) \ (\text{esp}) ]$ push buffer onto stack
[ 30   call #(readInputLoc) ] call the readInput function
[31   Ido  (1) #(tmpAddr) (0) ] Retrieve src address
[32   Ido  (2) #(destAddr) (0) ] Retrieve dest address
$\begin{bmatrix} 33 \\ 100 \\ 100 \end{bmatrix}$ = Retrieve the length
$\begin{bmatrix} 34 \\ \text{sto} \$(1) \#(4) \$(\text{esp}) \end{bmatrix}$ Push parameters on stack
$\begin{bmatrix} 35 &   sto s(2) \#(8) &   s(esp) \end{bmatrix}$
$\begin{bmatrix} 36 & 5(3) & 7(12) & 5(65) \end{bmatrix}$
$\begin{bmatrix} 3/  \operatorname{call} \#(\operatorname{strncpyLoc}) \end{bmatrix}$ call the string copy function
$\begin{bmatrix} 38 & 100 & (1) & f(srcAddr) & (0) \end{bmatrix}$ load the source
$\begin{bmatrix} 39 \\ 100 \\ 111 \\ 100 \end{bmatrix}$ #(destAddr) $(0) = 1$ = 10ad the dest address
$\begin{bmatrix} 40 & 100 & 5(3) & \#(100 & 100 & 100) & 100 &$
$\begin{bmatrix} 41 & \text{sto} & \text{s}(1) & \text{#}(4) & \text{s}(esp) \end{bmatrix}$ push the parameters
$[42   sto \mathfrak{F}(2) \#(0) \mathfrak{F}(esp) ]$
$[43   \text{stu} \phi(3) \# (12) \phi(\text{csp})]$ [ [44] call #(etrnempLec)] call structure function
$[44]$ (an $\pi(\operatorname{sum}(\operatorname{npLOC}))$ $[45]$ (begin $\pi(\operatorname{sum}(\operatorname{npLOC}))$ $[45]$ (begin $\pi(\operatorname{sum}(\operatorname{npLOC}))$ $[46]$ $(48)$ $[1]$ $(48)$ $[$
[45   000 46]
$[40]$ movi $\phi(1) \pi(0) = 1$ unequal $[47]$ begin $\xi(0) \#(0) \#(40) = 1$ on to the and
$[47]$ begin $\varphi(0) \#(0) \#(47)$ ] go to the end $[48]$ movi $\Re(1) \#(1)$ ] equal
[40 return] = 1

Figure 51: Assembly code corresponding to Figure 2

**Case 1:** Assume that the insider has corrupted the value of register \$3, which contains the length of the string. To model this attack, SymPLAID replaces the value in register \$3 with the symbol (*err*) just prior to calling the *strncmp* function. Consequently, the value

*err* is pushed onto the stack as an argument to the function. The *strncmp* function loads the value *err* into register \$(18) (instruction 2 of *strncmp*). It then initializes the result of the comparison in register \$1 to 1, and evaluates if the value in register \$18 is lesser than 0 (instruction 4 of *strncmp*). If so, it branches to instruction 15 of the *strncmp* function, which in turn returns the value in register \$1 to the *authenticate* function.

When SymPLAID encounters the comparison instruction 4 (in *strncmp*), it cannot uniquely resolve the comparison as it has no information on the value in register \$18 (which is *err*). Consequently, it forks the execution into two – one fork sets the value of register \$1 to be lesser than 0, and the other sets the value to be greater than or equal to 0. The former case takes the branch to instruction 15, and exits the function, thereby returning the value 1 to the *authenticate* function. This causes the *authenticate* function to return the value 1 to its caller – which is the desired goal of the attacker. The latter fork in which the value of register \$(18) is set to be greater than 0 does not however result in the outcome desired by the attacker. Therefore, in the above example, SymPLAID discovers that by setting the value of the length parameter of the *strncmp* function to a value lesser than 0, the attacker successfully achieves the goal of authenticating the user.

**Case 2:** Assume that the attacker overwrites the value in register \$2, which holds the address of the *dest* string. This is passed as an argument to the *strncmp* function, and copied into register \$17 by instruction 2 of the function. SymPLAID represents the value in register \$2 by the symbol *err* and tracks its propagation to register \$17. This value is used as a pointer argument to by the load-byte-unsigned (instruction 6) in *strncmp*. At this point, symPLAID cannot uniquely determine the memory address that the register

215

refers to, and hence it forks the execution so that each fork evaluates the symbolic expression to a different address.

The memory state of the program is generated by executing the code of the *authenticate* function until the fork point. SymPLAID assigns the value in register \$17 to each of the memory addresses in succession. Simultaneously, it assigns the value loaded into register \$6 to the value stored in the corresponding address. Of all the possibilities, only one of them leads to a state that satisfies the attacker's goal (which is that the *authenticate* function returns 1). This occurs when the address of the *src* buffer is passed to the function, as it will match the other argument which is also *src*.

**Case 3:** A similar case arises when the value in register \$1 is corrupted prior to the call to *strncmp*. The value in register \$1 holds the value of the string *src*, which is passed as the first argument to the *strncmp* function. The value is stored in register \$16 by instruction 1 and is used as as a pointer in instruction 9 of the strncmp function. As in the previous case, SymPLAID assigns a value of *err* to the register, and forks the execution at instruction 9, assigning each memory address in succession to the pointer value in register \$16. Of all the values considered by SymPLAID, only two values (*dest* and *tmp*) satisfy the attacker's goal, and are hence reported by SymPLAID as attacks.

**Summary:** Note that the output from SymPLAID is in the form of raw data and may consist of repetitive or redundant states. The output needs to be post-processed in order to identify more generic classes of potential attacks to drive the development of defense

216

mechanisms that can protect the application against a range of discovered attacks. The post-processing is currently done manually.

# 6.6 CASE STUDY

To evaluate the SymPLAID tool on a real application, we considered a reduced version of the OpenSSH application [137] involving only the user-authentication part. This is because SymPLAID does not support all the features used in the complete SSH application, e.g. system calls. We retain the core functions in the authentication part of OpenSSH with little or no modifications, and replace the more complex ones with stub versions – i.e. simplified functions that approximate the behavior of their original versions. We also replace the system calls with stubs. The reduced version is called the authentication module.

The authentication module emulates the behavior of the SSH application starting from the point after the user enters his/her username and password to the point that he/she is authenticated (or denied authentication) by the system. The authentication module consists of about 250 lines of C code (excluding standard libraries). The functions in the module are shown in Table 26.

Function Name	LOC(C)	Functionality
fakepw	15	Fills a structure with a default (fake) password and returns it
shadow_pw	7	Stub version of a system call to retrieve the hash of the password
getpwnam	19	Stub version of a system call to retrieve password for a username
рwcopy	22	Makes a field-by-field copy of the password structure
sys_auth_passwd	29	Checks if the user supplied password matches system password
allowed_user	6	Stub version of a complex function to check if a user is in the list of allowed users
xcrypt	7	Stub version of a system call to encrypt the password using a salt value (based on username)
getnyram allow	43	Checks if a user is allowed to login and if so retrieves their password record makes a copy
getpwnain-anow		using pwcopy
auth_password	14	Checks if the username is allowed AND the user password is correct
main	47	Reads in the username and password and calls the above functions in the expected order

Table 28: Functions in the OpenSSH authentication module

We ran SymPLAID on the authentication module after compiling it to MIPS assembly using the *gcc* compiler. As before, the goal is to find insider attacks that will allow the user to be authenticated. It is assumed that the insider can overwrite the value of a register in any instruction within the authentication module. The input to the authentication module is the username and password. The username may or may not be a valid username in the system, and the password may or may not be correct. These lead to four possible categories, of which one is legitimate and three are attacks. SymPLAID discovered attacks corresponding to the categories where an invalid username is supplied with a valid password (for the application) and where a valid user-name is supplied with an incorrect password. In this section, we consider both categories of attacks.

#### 6.6.1 Category 1: Invalid User-name

The authentication part of SSH works as follows<sup>23</sup>: when the user enters his/her name, the program first checks the user-name against a list of users who are allowed to log into the system. If the user is allowed to log into the system, the user record is assigned to a data-structure called an *authctxt* and the user details are stored into the *authctxt* structure. If the name is not found on the list, the record is assigned to a special data-structure in memory called as *fake. fake* is also an *authctxt* structure, except that it holds a dummy username and password. This ensures that there is no observable difference in the time it

 $<sup>^{23}</sup>$  We consider only the case where authentication is done using the keyboard.

takes to process legitimate and illegitimate users (which may enable attackers to learn if a username is valid by repeated attempts to login).

In order to prevent potential attackers from logging on by providing this dummy password, the *authctxt* structure has an additional field called *valid*. This field is set to *true* only for legitimate *authctxt* records i.e. those for which the username is in the list of valid users for the system. The *fake* structure has the *valid* field set to *false* by default. In order for the authentication to succeed, the encrypted value of the user password must match the (encrypted) system password, *and* the valid flag of the *authctxt* record must be set to the value 1.

Figure 53 shows the *auth\_password* function that performs the above checks. The function first calls the *sys\_auth\_passwd* to check if the passwords match, and then checks if the *valid* flag is set in the *authctxt* record. Only if both conditions are true will the function return 1 (authenticated) to its caller.

int	sys_auth_passwd(Authctxt *authctxt, const char *password) {
1:	struct passwd * <i>pw</i> = authctxt->pw;
2:	char *encrypted_password;
3:	char * <i>pw_password</i> = authctxt->valid ?
4:	shadow_pw(pw) : pw->pw_passwd;
5:	if $(strcmp(pw_password, "") == 0 \&\&$
6:	strcmp(password, "") == $0$ )
7:	return (1);
8:	encrypted_password = xcrypt(password,
9:	(pw_password[0] && pw_password[1]) ?
10	pw_password : "xx");
11:	return (strcmp(encrypted_password, pw_password) == 0);
}	
int	auth_password(Authctxt *authctxt, const char *password) {
12	int permit_empty_passwd = 0;
13	struct passwd * pw = authctxt->pw;
14	int result, ok = authctxt->valid;
15	if $(*password == '\0' \&\& permit_empty_passwd == 0)$
16	return 0;
17	result = sys_auth_passwd(authctxt, password);
18	if (authctxt->force_pwchange)
19	disable_forwarding();
20	return ( <i>result &amp;&amp; ok</i> );
}	

Figure 52: SSH code fragment corresponding to the attack

An insider can launch an attack by setting the *valid* flag to true for the *fake* authctxt structure. This will authenticate a user who enters an invalid user name, but enters the password stored in the *fake* structure. The password in the *fake* structure is a string that is hardcoded into the program.

To mimic this attack, we supply an invalid user-name and a password that matches the *fake* (dummy) password. We expected SymPLAID to find the attack where the insider overwrites the valid flag of the *fake* structure. SymPLAID found this attack, but it also found a host of other attacks that overwrite the frame pointer of the function. We describe a particularly interesting attack found by SymPLAID here.

The attack occurs in the *sys\_auth\_password* function, at line 11 before the call to the *strcmp* function (in Figure 52) .At this point, the insider corrupts the value of the stack pointer (stored in register \$30 in the MIPS architecture) to point within the stack frame of the caller function, namely *auth\_password*. When the *strcmp* function is called, it pushes the current frame pointer onto the stack, increments the stack pointer and sets its frame pointer to be equal to the value of the stack pointer (corrupted by the attacker). Figure 53 shows the stack layout when the function is called (only the variables relevant to the attack are shown).



Figure 53: Stack layout when strcmp is called

The top-row of Figure 53 shows the frame-pointers of the functions on the stack due to the attack. Observe that the attack causes the stack frame of the *strcmp* function to overlap with that of the *auth\_password* function. The *strcmp* function is invoked with the addresses of the *encrypted\_pasword* and the *pw\_password* buffers in registers<sup>24</sup> \$3 and \$4. The function copies the contents of these registers to locations within its stack frame at offsets of 4 and 8 respectively from its frame pointer. This overwrites the value of the local variable *ok* in the *auth\_password* function with a non-zero value (since both buffers are at non-zero addresses). When the strcmp returns, the value of \$30 is restored to the frame pointer of *sys\_auth\_passwd*, which in turn returns to the *auth\_password* function. The *auth\_password* function checks if the result returned from *sys\_auth\_password* is non-zero and if the *ok* flag is non-zero. Both conditions are satisfied, so it returns the value 1 to its caller, and the user is successfully authenticated by the system.

# 6.6.2 Category 2: Incorrect Password

The second category corresponds to the case when the application is executed with a valid username but with the wrong password. We ran SymPLAID on the application and asked it to find attacks where the user is successfully authenticated. We consider a particularly interesting example attack found by SymPLAID. The attack is described in this section.

<sup>&</sup>lt;sup>24</sup> In the MIPS architecture, function arguments are passed in registers

The attack occurs in the function *sys\_auth\_password* shown in Figure 52. As can be seen from the Figure, the function *sys\_auth\_password* returns 1 to its caller (*auth\_password*) if either the encrypted version of the user password matches with the encrypted version of the system password, OR if both passwords are empty strings. In a normal execution of the SSH application, the user password is checked by the *auth\_password* function, and if empty, a special flag *permit\_empty\_password* is checked. This flag indicates if the user is allowed to have an empty password (at account creation time, for example). If the flag is not set, the application is aborted. Therefore, under normal circumstances, the user password cannot be empty. However in the case where it is empty, the *auth\_password* function returns a value '1' provided the corresponding system password is also empty. A naïve attacker may try overwriting the value of *permit\_empty\_passwd* and entering an empty string for the password. However, this would require that the system password is also empty. Since we assume that only one corruption is allowed per execution, the attacker will not be able to corrupt both the system password and the user password simultaneously to make both of them point to empty strings, and the attack will not succeed. A better option for the attacker may be to overwrite the value of the system password (*pw\_password*) after it is returned from the *shadow\_pw()* function. This would not work either as the attacker would need to overwrite the user password (*authctxt->pw*) in order for the attack to succeed, which is not possible given the single value restriction. To craft a successful attack, observe that the system password is returned by the function shadow\_pw (since the username is valid, authctxt->valid is set to 1). Therefore, the attacker can try to make *shadow\_pw* return an empty string and in the process, also

222
overwrite the contents of the *password* variable. The difficulty with this approach is that *shadow\_pw* is a system call and its value is determined based on the value of the system password. Nonetheless, it is possible to make *shadow\_pw* return an empty string by passing it a NULL string as argument. This can be done by shifting the frame pointer of the *sys\_auth\_passwd* function to a memory location where the value stored in the address corresponding to the *pw* variable is 0, AND the value corresponding to the *password* variable points to an empty string. The attack is carried out after the check on *authctxt->valid* but before the call to *shadow\_pw* at line 4 in Figure 52.

## 6.6.3 Summary: Attacks Found

Table 29 summarizes the attacks discovered by SymPLAID for the two categories presented in sections 6.6.1 and 6.6.2. The attacks shown in Table 29 are confined to the two functions shown in Figure 52. We do not consider attacks that originate in the stubs for the system calls or those that originate in the main function as these are artifacts of the authentication module, rather than the application. We have validated the attacks shown in Table 29 on the OpenSSH application compiled for the MIPS architecture. We used the SimpleScalar emulator for carrying out the insider attacks.

FUNC	САТ	Attack found by SymPLAID for the function and category shown in the first two columns
auth_password	1	After initializing the value of local variable <i>ok</i> with <i>authctxt-&gt;valid</i> , overwrite it with a non-zero value.
	1	After setting up the stack with the local variables and return address, but before calling <i>sys_auth_passwd</i> , overwrite the stack pointer with an address such that <i>sys_auth_passwd</i> overwrites the value of the variable <i>ok</i> when storing on the stack
	2	After the call to sys_auth_password, overwrite its return value (register \$2) with a non-zero value
	2	After the call to <i>sys_auth_password</i> , shift the frame pointer of the function such that before the function returns, it will read non-zero values in the relative addresses of the local variables <i>result</i> and <i>ok</i>
	1, 2	Overwrite the value of register \$2 with a non-zero value when the function is about to return, so that it always returns <i>true</i>
auth_password	1	Before the call to <i>strcmp</i> , inject the stack pointer with an address in the previous stack frame so that the value of <i>ok</i> is overwritten with a non-zero value. This makes <i>auth_password</i> return 1.
	1	Before the call to the strcmp function, overwrite the frame pointer with an address in the previous stack frame, so that when <i>strcmp</i> returns a non-zero value (since the strings are different), the return address is the address of the instruction that called <i>auth_password</i> (from <i>main</i> ).
	2	After the call to the <i>xcrypt</i> function, set the value returned by it in register \$2 to the address of the buffer where the encrypted system password is stored. This sets the user password pointer to the encrypted value of the system password.
	2	Before calling the <i>strcmp</i> function, change either its first argument (register \$4) to the address of buffer with the encrypted system password, or change the second argument (register \$5) to the address of the buffer containing the encrypted user password
	2	After the call to the shadow_pw address, change its return value to the address of the buffer that contains the user password
sys_	2	Change the return value of the function (register \$2) to a non-zero value

Table 29: Summary of attacks found by SymPLAID for the module

# 6.6.4 Spurious Attacks

The approximations in the authentication module may introduce spurious attacks. These are attacks that work on the authentication module but do not work on the real OpenSSH application. This is because the stub functions in the module introduced approximations that did not mimic the real system's behavior in all cases. Since the model-checker explores all possible behaviors of the system, it flagged the non-conforming cases as attacks. A similar phenomenon was observed by Musuvati and Engler [139].

Most of the spurious attacks discovered were easy to filter out as they were launched

from stub functions that were system calls in the real program. However, there was one

subtle attack that at first seemed like a real attack, but turned out to be spurious. This is described here in this section.



Figure 54: Schematic diagram of chunk allocator

The OpenSSH program uses its own custom memory allocator (*xmalloc*) to store the buffers containing the user and system passwords (not shown in Figure 52). *xmalloc* allocates memory in chunks, and if it runs out of space in the chunk, it calls the system malloc to allocate a new chunk. Our authentication module simplifies this behavior by allocating a single static chunk of memory when the application is initialized, and then satisfying all application malloc requests from this chunk. The initial chunk is chosen to be large enough to accommodate both the password buffers and other dynamic memory used in the program. In order to ensure that we do not exceed the size of the initial chunk, every allocation request in the program is checked to ensure that it is within the bounds of the space remaining in the chunk.. Figure 54 shows a schematic of the allocator.

The simplified memory allocator has a runtime check of the form:

*if* (*currentPos* + *sizeRequested* > *maxSize*) *return NULL*;

where *currentPos* is the location of the next free location in the initial chunk, and *maxSize* is the size of the initial chunk.

The SymPLAID tool finds an attack that effectively overwrites the location of the *currentPos* variable in memory with a value that is greater than the value of *maxSize*. This causes all subsequent *malloc* requests from the application to be declined and the value NULL to be returned. In order to prevent a NULL pointer violation, the calling function makes the pointers point to a special sentinel value in memory. If this attack is carried out at the beginning of the authentication module (in the *getpwnamallow* function, say), this will cause both the password strings to point to the same sentinel value, and hence they will match with each other. Therefore, the *sys\_auth\_password* function will return 1, and the malicious user will be authenticated, thereby leading to a successful attack.

In reality, this attack cannot be achieved easily as the custom allocator in *ssh* will not merely return NULL if it exceeds the bounds of the current chunk, but will get a new chunk from the operating system (using *brk* in linux), and maintain a linked list of the allocated and free chunks. It is conceivable that a more sophisticated version of the attack can be mounted by overwriting the head of the free list with NULL to simulate the conditions leading to memory exhaustion. However, we have not tested the more sophisticated attack.

The above situation could have been avoided had we modeled a more accurate version of the memory allocator used by OpenSSH. However, a similar situation could have arisen in any of the other stub functions. Therefore, any approximation of system behavior is likely to lead to spurious outcomes. The only way to avoid this situation is to analyze the entire system (application, libraries and operating system) using the model-checker as

done in [139]. However, this can lead to state space explosion in the model-checker and is an area of ongoing work.

#### 6.6.5 Performance Results

This section reports the performance overheads incurred by SymPLAID in finding the attacks on the OpenSSH application. We executed SymPLAID on a parallel cluster consisting of dual-processor AMD Opteron nodes, each of which has 2 GB RAM. This is because the search task is highly parallelizable and can be broken into independent sub-tasks, with each sub-task considering attacks in a different code region of the application. The authentication module consists of about 500 assembly language instructions, and the task was broken up into 50 parallel sub-tasks, each of which considers a code region of 10 instructions. The maximum time allowed for completion of a sub-task is approximately 2 days (after which the task is forcibly terminated) and its execution time recorded as 48 hours.

Function Name	LOC	Number of States	Total Time (sec)	Attacks found
	(assembly)			?
getpwnamAllow	37	6391	325861	No
sys-auth-passwd	54	36896	366108	Yes
fakepw	29	11	115	No
xcrypt	37	26921	429683	Yes
shadow-pw	26	10342	272236	Yes
allowed-user	20	11	115	Yes
auth-password	40	26921	429683	Yes
getpwnam	37	27724	534601	Yes
pwCopy	52	23547	471185	Yes
main	114	3137	297526	Yes

Table 30: Time taken by SymPLAID for each function

Table 30 shows the time and space requirements of the sub tasks categorized by the function which they were analyzing for attacks. The space requirements are reported in

terms of the number of unique "states" visited by the model-checker. The time taken is reported in seconds. The results are aggregated across multiple sub-tasks for the function and the cumulative time and space requirements are reported. Note that this is not equivalent to running the sub-tasks for the function as a single aggregate task as the sub-tasks may have significant state sharing across them. Hence the time and space taken by a single aggregate task is likely to be smaller than the aggregated results in Table 30. Based on the results in Table 30, the total time taken to execute all sub-tasks is at most 3127113 seconds or 36.2 days. However, we are able to finish the task in 2 days due to the highly parallel nature of the search task. The total number of states explored by the sub-tasks is 161091.

While the running time seems high, it is not a concern as the goal is to discover all potential attacks (in a reasonable time frame) and to find protection mechanisms against them. The analysis can be easily parallelized and executed on multiple nodes as independent sub-tasks (as we did). Therefore, as we move towards multi-core and large-scale parallel computers, the analysis time is bound to decrease. Finally, model-checking is a very active area of research and new techniques are being invented to make model-checking faster. We can take advantage of such approaches to reduce the running time of SymPLAID.

The reasons for high running time are: (1) SymPLAID performs an exhaustive search of memory locations whenever it encounters a load/store through a register containing a symbolic expression., (2) When SymPLAID encounters an indirect jump instruction with a register containing a symbolic expression, it needs to scan the entire code-base as

potential targets for the jump, and (3) SymPLAID distinguishes between states with minor differences in their representation (not relevant to attack).

We are currently investigating the following optimizations to reduce the running time.

- Instead of performing an exhaustive search of memory locations, consider multiple locations in the form of abstract memory regions. Further, if memory safety-checking techniques are deployed, it is enough to consider memory accesses within the write-set of a pointer location.
- Restrict the scope of indirect jumps to be within a function or module. Controlflow checking techniques to reduce the scope of valid jump targets in a program can be deployed.

#### 6.7 RELATED WORK

We classify related work into three broad categories as follows: (1) Identification and generation of insider attacks, (2) Symbolic execution techniques to find security vulnerabilities and, (3) Fault Injection techniques to perturb application state.

#### 6.7.1 Insider Attacks

Insider attacks have been a significant source of security threats, and efforts have been made to model insider attacks at the network level. Philips and Swiler[134] introduced the attack graph model to represent the set of all possible attacks that can be launched in a network. The nodes of an attack graph represent the state of the network, and each path in the graph represents a possible attack.. Ritchey and Amman [133] introduce a model-

checking based technique to automatically find attacks starting from a known goal state of the attacker. Sheyner et. al. generalize this technique to generate all possible attack paths, thereby generating the entire attack graph [132]. Chinchani et al. present a variant of attack graphs called key-challenge graphs that are specifically tuned to represent insider attacks [140].

Insider attacks were also modeled at the operating system level by Probst et al.[135]. In this model, applications are represented as sets of processes that can access sets of resources in the system. An insider is modeled as a malicious process in the system that may access resources in violation of the system's security policy. Their technique builds a process interaction graph for the system and performs reachability analysis to discover insider attacks.

Attack-graphs and process graphs are too coarse grained for representing applicationlevel attacks, and hence we directly analyze the application's code. Further, we do not require the developer to provide a formal description of the system being analyzed, which can require significant effort. Since we analyze the application's code directly, we can model attacks both in the design and implementation of the application. This is important as an insider typically has access to the application's code, and can launch low-level attacks on its implementation.

## 6.7.2 Symbolic Execution

Symbolic execution is a well-explored technique to find program errors [117]. Recently, it has also been used to find security vulnerabilities in applications [141-144].

Kruegel et al. present a technique to automatically generate mimicry attacks against system-call based attack detection techniques [143]. By symbolically executing the program, their technique can find attack inputs that can execute a malicious system call, while replicating its context, thereby remaining undetected (by the monitor).

EXE is a symbolic execution technique to generate security attacks against applications [144]. Their approach directly executes the application code on symbolic inputs, and progressively constrains them when conditional branches or assertions are encountered. Molnar and Wagner generate attacks to exploit integer conversion errors in programs [142]. Their technique starts with a valid (non-attack) input and attempts to mutate it into an input that exploits a given integer conversion vulnerability.

Bouncer generates filters (program assertions) to block exploits of known memory corruption vulnerabilities [141]. The technique starts with an attack that exploits a given vulnerability, and symbolically executes the program to generate a set of constraints under which the vulnerability can be exploited. It then uses the generated constraints to block all inputs that may exploit the vulnerability.

The above techniques are concerned with generating attack inputs for the applications to exploit known or unknown vulnerabilities. In contrast, our technique attempts to generate attacks for a given input, assuming that the attacker is already present in the system. Further, the attacks found using our technique do not require the application to have an exploitable vulnerability (e.g. buffer overflows), but can be launched by a malicious insider in the system.

#### 6.7.3 Fault-Injection Techniques

Fault-injection is an experimental technique to assess the vulnerability of computer systems to random events or faults [19]. An artificial perturbation signifying a random fault is introduced in the system and the behavior of the system is studied under the perturbation. Traditional fault-injection is statistical in nature, and is hence not guaranteed to expose all corner scenarios in the application. Consequently, it is not wellsuited for modeling security attacks, as attackers typically exploit corner-case or unexpected behaviors. In spite of these limitations, researchers have used fault-injection to find security violations in systems. We consider some examples of such techniques. Boneh, DeMillo and Lipton pioneered a study in which they found that transient hardware errors could adversely affect the security guarantees provided by public-key cryptosystems [145]. Subsequent studies have shown that many commonly used cryptographic systems can be broken by hardware errors in their implementation [146]. The main difference between these studies and ours is that our technique can be applied for any general security-critical system rather than only crypto-systems. Further, we allow the attacker to inject any value into the processor's registers or memory which are more illustrative of insider attacks.

Xu et al. consider the effect of transient errors (Single-bit flips) in control-flow instructions on application security [147]. They use a technique known as "selectiveexhaustive" injection to inject all possible errors in code segments that are known to be critical to the integrity of the systems from a security point of view. Our technique may

be viewed as a form of selective-exhaustive injection, but into program data rather than instructions. The other difference is that we consider the effect of all possible data value corruptions rather than just single bit flips.

Govindavajhala and Appel [148] explore the use of transient errors to attack a virtual machine when the attacker has physical access to the machine. They show that transient errors can break the protections of the virtual machine up to 70% of the time, depending on the platform and the attacker's ability to execute a specially crafted program. The main difference between this work and ours is that we do not require attackers to have physical access to the machine, nor launch specially crafted applications.

## 6.8 CONCLUSION

This chapter presented a novel approach to discover insider attacks in applications. An automated technique to find all possible insider attacks on application code is presented. The technique uses a combination of symbolic execution and model-checking to systematically enumerate insider attacks for a given goal of the attacker. We have implemented the technique in the SymPLAID tool, and demonstrate it using the code segments corresponding to the authentication part of the OpenSSH application. We find several instances of potential insider attacks which may be missed by simple, manual inspection of the code.

# CHAPTER 7 INSIDER ATTACK DETECTION BY INFORMATION-FLOW SIGNATURE (IFS) ENFORCEMENT

# 7.1 INTRODUCTION

The growing complexity of applications has necessitated a shift towards outsourcing of application components to third-party vendors often spanning geographic boundaries. The ubiquity of the internet has allowed software libraries to be freely distributed in source/binary forms and reused among applications. In this environment, a malicious developer may plant a logical loophole or backdoor in a library or module used by a security-critical application. The developer could leak details about the loophole to an attacker, who could then exploit the application through the loophole or backdoor when it is deployed in the field. Such backdoors and loopholes are extremely hard to detect unless detailed code auditing is performed (if the source code of the module is available). However, due competitive pressures in bringing a product to market, organizations often do not perform these tasks for code that is not developed in-house. Even for modules developed in-house (i.e. within the organization), the original developer(s) may have long left the company and the expertise to understand the code may be lost.

This problem is partially alleviated by open-source software, due to the "many eyeballs" scanning the source code for potential loopholes[149]. However, even in open-source software, it is possible for a malicious developer to plant a backdoor or loophole in an unused (or rarely used) part of the code. While backdoors have been ferreted out in popular open-source packages such as the Linux Kernel [150], they may not be as easy to

detect in open-source packages that are not as widely-deployed. For example, a study of open-source software packages distributed using the popular SourceForge repository found that about 50 % of the packages had fewer than 70 downloads during the month-long duration of the study [151]. Further according to [152], a recent study of 100 open-and closed-source software packages found that about 23 of them had unwanted code, and about 79 packages had dead code of some form or another, although they were not necessarily malicious.

It is also possible for a malicious system administrator or IT manager to replace a system library or in-house software package with a modified version, in order to subvert existing security checks or induce malicious behavior. A recent CERT report on insider attacks [153] shows that while such attacks are relatively rare (only 15 out of 200 cases studied in the report belonged to this category), the attacker can cause extensive damage through these means. For example, the report illustrates how an IT manager in a state agency was able to carry out an extensive fraud-scheme undetected for two years by commenting out a single-line of source code in an in-house software program, and compiling and releasing the modified version within the organization.

Malware is defined as any program that attempts to perform malicious activities in the system [154]. Malware takes many forms, including viruses, worms and Trojans. Viruses and worms attach themselves to executable files and are activated when the program is executed. Trojans are stand-alone programs that masquerade as programs with legitimate functionality in order to trick users into executing the Trojan program.

Both static and dynamic approaches have been proposed for malware detection. Static malware detectors (including commercial virus scanners) look for known patterns of instructions representing a virus or Trojan in executables. However, such static checkers can be bypassed by malware that performs simple program obfuscations on itself [155]. Christodorescu et. al. [156] takes into account instruction semantics to determine if a given instruction sequence is a semantic variant of a known malware sequence. This approach requires the original sequence template of instructions in the malware to be specified, and variants are automatically detected by the technique. Dynamic approaches for malware detection check whether a program (i.e. potential malware) performs system-level malicious activities such as overwriting operating system files or registry entries. The check can be performed by monitoring the program using software [157] or hardware [158]. These approaches attempt to emulate conditions in the field in order to trick the malware into revealing its malicious behavior.

In contrast to malware, an untrusted module in an application does not attempt to infiltrate the system by performing system-level malicious activities. Rather the module may overwrite key elements of the application's control and data-space to achieve the attacker's goals. Hence, dynamic malware detection approaches will not be effective at detecting the attack. Further the instruction sequences corresponding to the malicious activities carried out by the module are specific to the application being attacked. Hence, the sequences may not be detected by a generic pattern-matching or semantics-based approach for static malware detection.

The attack model assumed in this chapter is that an untrusted module (whose source code may or may not be available) has been linked together with the application prior to the application's deployment. The untrusted module has a malicious snippet of code hidden in it (possibly masquerading as legitimate code) which performs some malicious activity in the context of the application e.g. overwrite security-critical data or bypass security checks. We refer to the untrusted module as an "application-level insider<sup>25</sup>". For example, an insider in a login program may overwrite the program's internal state to allow an attacker to log into the system even if he/she does not provide the correct password.

Analyzing an application's source code for application-level insiders is a hard problem. This is because the malicious code can masquerade as code with some legitimate functionality and pass manual audits[152]. Automated code analysis techniques can detect the malicious code segment, but these techniques require a specification of the program to be provided. In order to detect insider attacks, every line of an application's source code must be checked against its specification (as it is a potential hiding place for malicious code). Formulating a specification for each line of code is cumbersome and few developers choose to write such detailed specifications. The problem is exacerbated at the binary level, as binary code is often stripped of source symbols and obfuscated precisely to inhibit its understanding by a tool or human. Further, the malicious code

 $<sup>^{25}</sup>$  An insider attack is defined as one in which a privileged entity i.e. an insider, abuses its privileges to exploit system loopholes or perform malicious activities. In this case, the privilege afforded to the untrusted module is that it is allowed to execute in the same address space as the application.

may not even be revealed by testing techniques as it may be activated only under a specific combination of inputs or environment variables. While it is possible to exhaustively test the program under every legal combination of inputs and examine its behavior, such exhaustive testing is not done due to practicality reasons of time and resource overheads.

The goal of this chapter is to develop a technique to detect application-level insider attacks. We focus on protecting a subset of the application's data that is critical to the security of the application from updates by untrusted modules (i.e. whose source code is neither available nor inspected). This *critical data* has to be identified based on the application's semantics either manually or automatically using specifications about the program.

The technique presented in the chapter uses the idea of Information-Flow Signatures (IFS) to detect application-level insider attacks that illegitimately overwrite critical data in the application. The programmer identifies the security-critical data through annotations in the source code. The IFS encodes the sequence of instructions that can legitimately write (directly or indirectly) to the critical data through the normal (attack-free) control-flow of the program. The IFS is derived using static analysis of the program's code (by enhancing the compiler), and requires the source-code of only those modules that can legitimately write to the critical data, i.e. the trusted modules. The program is instrumented to ensure that the statically-derived IFS is followed at runtime – any deviation represents an attack and the program is halted. Note that only attacks that violate the integrity of critical data are detected by the IFS. We do not consider attacks on

the confidentiality of the critical data or Denial-of-Service (DoS) attacks on the application.

In earlier work [159], we introduced the idea of Information-flow signatures (IFS). This chapter builds significantly on the initial idea and shows how to derive and enforce the IFS technique for real applications. The derivation of the IFS is implemented by enhancing the LLVM optimizing compiler [99], while the runtime enforcement is done via custom software libraries. The technique is demonstrated in the context of three commonly deployed server applications (OpenSSH, WuFTP and NullHTTP) to protect security-critical modules. *We find that the performance overheads of the proposed technique range from 7.5 % to 100 % when measured with respect to the execution time of the module(s) protected, but are negligible (less than 1 %) when evaluated in the context of the entire application as the protected module constitutes only a small part of the application's execution time.* 

The security guarantees provided by the IFS technique depend on the choice of critical data in an application. We demonstrate in [160] an automatic technique to choose critical data by systematically enumerating all possible attacks that may be launched by an insider. In order to deploy the technique in [160], the user has to provide generic specifications about the attacker's goal (for example, to log in with the wrong password) and the system will automatically identify the critical variables in the application to be protected in order to foil the attacker's goals). However, for the applications considered in this chapter, the critical data was chosen manually based on our understanding of the

application's source  $code^{26}$ . We show experimentally that the technique can detect both insider attacks and external attacks that impact the critical data in each application.

The chapter makes the following contributions:

- Proposes a novel technique to protect critical data in applications from insider attacks through the concept of Information-Flow Signatures (IFS)
- Leverages and enhances existing static analysis techniques to extract the backward slices of critical variables and instrument the instructions in the slice to derive the IFS of the application (Section 4).
- Analyzes the efficacy of the proposed technique against both generic and targeted attacks (Section 5). Also, discusses the effect of approximations (made in the prototype) on the technique's effectiveness.
- Evaluates the performance overheads of the IFS technique for the benchmark applications and the resilience of the applications protected with the IFS technique to attacks (Section 7).
- Proves the efficacy of the IFS technique in detecting insider attacks, and shows that the IFS technique detects all external attacks impacting critical data that are also detected by existing techniques (see Section 7.10).

<sup>&</sup>lt;sup>26</sup> This is because the formal technique could not be scaled to the applications presented in this chapter. This has to do with the limitations of the model-checker used in the earlier work, and is orthogonal to the technique for identifying critical data.

# 7.2 RELATED WORK

Insider attacks are attacks in which a privileged entitiy ("insider") abuses its privilege to attack the system. Insider attacks can be mounted at the hardware, virtual machine, operating-system, and network levels. Table 31 provides a brief overview of techniques deployed at each level to foil insider attacks. Each technique ensures that the insider cannot infiltrate the system at the level in which the technique is deployed. However, none of these techniques address the problem of application-level insiders. This is because an application-level insider does not need to change the hardware, operating system, virtual machine or network in order to launch its attack<sup>27</sup>.

Layer	Techniques	Comments
Hardware	King et al. [161]	These techniques consider malicious backdoors in hardware either at the design
	Alkabani et. al.	stage or at the synthesis stage. At the design stage [161], the problem is similar to
	[162]	the application-level insider attack problem (and there is no known solution). At
		the synthesis stage [162], the problem is similar to tampering with the
		application's executable file and can be alleviated by embedding hidden keys in
		the synthesized netlist for comparison with the original netlist.
Virtual Machine	King et al [163],	The insider controls the boot-process prior to the loading of the Operating System
	Rutkowska	(OS) and Virtual Machine (VM) and loads their own malicious VM in order to
	[164]	host the OS/VM [163]. From the malicious VM, the insider can wreak
		considerable damage totally transparent to the operating system or application.
		Later work has shown that it is possible to launch the attack even after the OS has
		finished loading (through hardware virtualization hooks) [164]. The malicious
		VM can be detected through timing measurements made from the guest VM/OS
		[165]
Network	Upadhyaya et al.	The attacker is assumed to control one or more nodes in the network through
	[166], Sheyner	which it is assumed that the attacker tries to launch attacks on other nodes. A
	et al., Philip and	structure called the attack graph is used to represent the possible malicious
	Swiler [132,	actions of the attacker [132, 134]. These attacks can be detected by enhancing
	134]	Intrusion Detection Systems (IDS) to look for anomalous patterns of traffic
		within the network [166] that are representative of a network-level insider.
Operating System	Rootkit	Rootkit detection looks for anomalous data-structures or memory access patterns
	detection	of the OS to determine if a malicious module has hooked into the OS (through an
	[167, 168]	existing vulnerability). These techniques are based on knowledge of the operating
		systems' state (or some approximation of it) prior to the infection by the root-kit.
	Micro-kernel	Structured the operating system into multiple, non-overlapping services. Each
	based OS [169,	service runs in its own privileged compartment, and communicates with other
	170]	services through a thin layer called the micro-kernel. So even if one service is
		compromised, the other services are not.

<sup>&</sup>lt;sup>27</sup> This is analogous to how reliability techniques at lower-layers of the system stack does not obviate the need for application-level protection. This is because errors that occur in the application are not detected by lower-level techniques such as ECC in memory.

The rest of this section discusses security techniques deployed at the application level, which protect the application from malicious attackers. We classify techniques for protecting applications from security attacks into three broad categories: (1) techniques to protect against external attackers (e.g. memory safety-checking, taintedness detection, address-space/instruction-set randomization and system-call based checking), and (2) techniques to protect against internal attackers including application-level insiders (e.g. privilege separation, remote audit, code attestation and oblivious hashing), and (3) techniques to protect critical data in applications from corruption due to errors and security attacks such as *Samurai* and *redundant data diversity*.

External security techniques such as memory safety checking [23, 171] are designed to protect applications from memory-corruption attacks (e.g. buffer-overflow, format string). This class of techniques is not effective for thwarting insider attacks as insiders do not need to exploit memory-corruption vulnerabilities in the application. Taintedness detection techniques [25-27] prevent application input from influencing high-integrity data in applications such as pointers. However, insiders do not need to use inputs to influence security-critical data as they are within the application. Further, it is almost impossible to prevent an internal module of the application from writing to generic program objects such as pointers, without incurring a very high-false positive rate. Security techniques such as randomization [28, 172] attempt to obscure a program's layout or instruction set from attackers. However, randomization can be bypassed by an insider who is itself subject to the same randomization as the application. For example, a malicious module in an application that is subject to address-space randomization (ASR)

can calculate the absolute address of a stack variable in a different function by examining the addresses of its local variables, and adding a fixed offset to them. Unlike in the case of an external attacker who requires multiple attempts to bypass the randomization [31, 173], an insider can bypass it in a single attempt.

System-call based checking is a technique that monitors the sequence of system calls made by an application and checks if the sequence corresponds to an allowable sequence as determined by static analysis techniques [174, 175]. These techniques assume that the attacker seizes control of the application by executing unwanted or malicious system calls such as *exec*, or by skipping existing system calls, e.g. *seteuid*. This is because system-calls provide a conduit to attack other applications executing on the same system as well as the Operating System (OS) itself. However, an application-level insider's goal is to subvert the execution of the attacked application, and not necessarily attack other applications or the OS. Hence, system-call based detection techniques will not protect against insiders who overwrite the attacked application's data or control without launching system calls.

Techniques such as oblivious hashing [176, 177] and code attestation [178] detect malicious modifications of the application's executable code after it has been generated (by the linker). However, these techniques do not protect from insider attacks in which the application developer links the application with an untrusted third-party library prior its distribution (as considered in this chapter). A technique that offers limited protection from insider attacks is remote audit [179], which ensures that the application's code is not

skipped at runtime. However, remote audit does not protect against malicious modifications of the application's data by an application-level insider.

The only known technique that can effectively thwart application-level insider attacks is privilege separation [180]. In this technique, the application is divided into separate processes and each module executes in its own process. A module can share data with another module only through the OS's Inter-Process Communication (IPC) mechanisms. This prevents an untrusted module from overwriting data in a trusted module's address space, unless the trusted module explicitly shared the data with the untrusted module (through an IPC call). However, privilege separation incurs overheads of up to 50 % when deployed in real applications [181] (measured as a fraction of the entire application's execution time, not just the protected module's execution time). Further, a trusted module may load an untrusted library function in its address space, thereby annulling the technique's security guarantees.

Finally, Samurai [182] and redundant data diversity [183] also protect critical data in an application from accidental and malicious corruption respectively. However, they both require the programmer to manually identify read/write operations on the critical data, which can be cumbersome. Further, Samurai only protects against corruption of critical data on the heap, and not for critical data on the stack or registers. Redundant data diversity requires replicating the entire process and executing it in lock-step, even though the critical data may constitute only a small portion of the application's data. This leads to unnecessary overheads and wasted resources.

Thus, there exists no technique that can detect insider attacks on program data without requiring considerable intervention on the part of the programmer or incurring high performance overheads. *The question we ask in the chapter is "Can we protect the integrity of critical data from insider attacks with low performance overheads and minimal intervention from the programmer?"* 

# 7.3 ATTACK MODEL

While the focus of this chapter is on insider attacks, the attack model also considers external attacks on the critical data of the application. This is because in addition to overwriting the critical data by itself, an insider can also plant a memory corruption vulnerability in the application which will be exploited by an external attacker to overwrite the security critical data. Hence, it is important to consider both external attacks and insider attacks on the critical data, as insider attacks are a super-set of external memory-corruption attacks on the application.

In the case of external attacks, we assume that the attacker exploits a memory corruption vulnerability (e.g., buffer overflow, format string) to overwrite critical data either directly or indirectly. A direct over-write means that the compromised instruction overwrites the security critical data through a pointer. An indirect over-write means that the compromised instruction overwrites a data element that influences the critical data through a data- or control- dependence. In both cases, the net effect is to influence the value of the critical data in such a way so as to benefit the attacker. The attacker may also

launch system calls on the program's behalf to impact the critical data<sup>28</sup>. However, the attacker is prevented from launching new processes through the system call interface, as such attacks would be detected by OS-level checking mechanisms.

In the case of insider attacks, we assume that an attacker can corrupt any program data or change the program's control-flow during calls to untrusted code that is controlled by the attacker. From within the untrusted code, the attacker can modify the value of any location in the stack, heap or processor registers in order to influence the critical variable. The attacker can also modify the return address or a function frame pointer on the application stack to force the program to return to a different address than the intended one. Note that an insider attack may be possible even if the application does not have any memory corruption vulnerabilities.

We do not assume that the source code of un-trusted third-party functions is available for analysis – however, it is assumed that the source code of all trusted modules is available. We also assume that the attacker cannot modify the application's code once it is loaded into memory. This is reasonable as in many systems the code segment is marked readonly after the program is loaded (unless the program is self-modifying).

# 7.4 APPROACH AND ALGORITHM

This chapter focuses on protecting the integrity of critical data from application-level insider attacks. Critical data is defined as any variable or memory object which, if

<sup>&</sup>lt;sup>28</sup> We assume that the attacker cannot infiltrate the Operating System (OS) by executing the system call and exploiting an OS vulnerability.

corrupted by an attacker can lead to security compromise of the application. In the proposed approach, the portion of the program that manipulates the critical data (i.e. the trusted module) is statically analyzed and instrumented with code to ensure that runtime modifications of the critical data follow the language-level semantics of the application. This corresponds to statically extracting the backward slice of the critical data, and ensuring that only the instructions within the slice can modify the critical data, and only in accordance with their execution order as specified by the program code. A third-party module or a memory corruption attack that overwrites the critical data violates the established dependencies in the slice and hence, can be detected. The approach ensures that information-flow to the critical data is in accordance with the program's source code, hence the name *Information-Flow Signatures (IFS)*.

**Invariants:** The instrumentation added by the IFS technique ensures that the following invariants are maintained.

- Only the instructions that are allowed to write to data operands in the backward slice of the critical data (according to the static data dependencies), in fact do so at runtime.
- 2. The instructions in the backward slices of the critical data are executed in the order of their occurrence along a specific set of acyclic paths in the program.
- 3. Either all the instructions in the backward slice along a specific path are executed at runtime, or no instruction along the path is executed.

**Overall Algorithm:** The algorithm for deriving and checking the IFS is split into four phases as follows:

#### Phase 0: Identification of Critical Data (Carried out by the programmer)

The critical data in a program can refer to both program variables (i.e. local and global variables) as well as dynamically allocated memory objects on the heap. The programmer identifies critical data in the program through annotations in the source code. In the case of program variables (local or global), the annotations are placed on the definitions of the variables<sup>29</sup>. In the case of memory objects, the annotations are placed on the allocation sites in the program (i.e. calls to *malloc*).

In this chapter, we use the term critical variable to refer to both critical variables and memory objects.

## Phase 1: Static Analysis: (Carried out by our enhancements to the compiler)

- Extract intra-procedural backward slice of the critical data by identifying all instructions within a function in the program that can influence the critical data. It is assumed that the function that manipulates the critical data is trusted, and its source code is available for the analysis.
- 2. For each instruction in the backward slice, insert an encoding operation after the instruction and pass the value computed by the instruction as an argument to the encoding operation.

<sup>&</sup>lt;sup>29</sup> We consider programs translated to Static Single Assignment (SSA) form, so each variable has a unique definition in the program.

- 3. Replace all uses of the instruction with the value returned by the encode operation within the function.
- Before every use of an operand that has been encoded in the backward slice, insert a call to the decoding operation and pass it the value returned by the encoding operation.
- 5. Generate the sequences of instructions for each function for each acyclic controlflow path in the function.
- 6. Add instrumentation functions at the beginning and end of function calls in order to push and pop the current state of the state machine on to a stack (see below).

## Phase 2: Code Generation: (Carried out by custom programs)

- 1. Generate finite-state machines to encode the sequences of calls to the encoding operations within a function for all control paths identified in step 5. Mark the final state of each state machine as an accepting state.
- 2. Generate the encoding and decoding operations to check the data-values of the program as it executes (see below). Also, check the validity of the program's control-flow using the state machines derived above in step 1.

## Phase 3: Runtime: (Carried out by the generated code)

 Track the runtime path based on the state machines generated in step 1. If the path does not correspond to a valid path, raise an alarm and stop the program (see explanation below). Encode data-values depending on where the encoding operation is called in the original program. Check the value for consistency by decoding it before its use, i.e. check if the decoded value matches the encoded operand. If the value is inconsistent, raise an alarm and stop the program.

**Slicing Algorithm:** The backward slices of the critical data are computed on a pathspecific basis, i.e., each execution path in the function is considered separately for slice extraction. This is based on our earlier work on extracting backward slices for detecting transient errors in programs<sup>30</sup> [184].

**Encoding/Decoding Operations:** The encoding/decoding operations protect the data in the backward slice after it is produced (enforce invariant 1). The encoding operation used in this chapter is duplication, where the operand is stored in a special, protected memory location. During decoding, the original value is compared with the stored value of the operand. A mismatch indicates that the operand has been tampered with, i.e., an attack. We assume that the attacker cannot modify the values stored by the encoding operations (as they are stored in protected memory).

It is possible to incorporate more advanced encoding functions such as checksums or even encryption to provide stronger protection, constrained by the incurred performance overheads. The above algorithm is orthogonal to the mechanism for protecting encoded operands.

<sup>&</sup>lt;sup>30</sup> The backward slice was used to *recompute* the value of selected program variables to check if they have been corrupted by an error.

**State Machines:** The state machines track the sequence of calls to the encoding functions and check if the program follows valid control-flow. The state is tracked on a per-function basis, since we consider only intra-procedural slices. A separate stack is maintained at runtime to push and pop the current state of the state machine (for the function) at the beginning and end of function calls. At the entry point to a function, the corresponding state machine is reset to the start state. Similarly, the state-machine's state is checked just before the function returns to ensure that the state machine is in an *accepting state* i.e., the state machine has accepted the observed sequence of encoding calls (invariant 3). Finally, we check that every encoding call executed by the program corresponds to a valid state transition from the current state of the state machine (invariant 2).

## 7.5 EXAMPLE CODE AND ATTACKS

This section illustrates the IFS technique using a code fragment drawn from the OpenSSH application. The section also considers example attacks on the application's code and discusses how IFS detects the attacks.

Consider the *sys\_auth\_password* function shown in Figure 55(a). The function accepts an *authctxt* data-structure and a password variable, and checks if the user password matches the password stored in the *authctxt* structure. If the passwords match, the function returns the value 1 and the user is authenticated by the system (not shown). The function also encrypts the user password prior to comparing it with the system password.

**Critical Data:** In order to determine the critical data, we assume that the goal of the attacker is to subvert the authentication mechanism by making the function return 1 in spite of the invalid user-name or password being given. Therefore, we designate the value returned by the function as the critical variable (i.e., the *authenticated* variable defined in line 6). The value returned in line 3 is not considered critical as it is a compile-time

constant.

0: int sys_auth_passwd(Authctxt* authctxt, const char*	Function Name	Purpose	Trusted ?
password) {	shadow_pw	Retrieves the shadow	Yes
1: struct password* pw = authctxt->pw;	<u>^</u>	password from the system	
2: char* pw_password = (authctxt->valid) ?		password file	
<pre>shadow_pw(pw) : pw-&gt;pw_passwd;</pre>	xcrypt	Computes an encypted	Yes
3: if (! strcmp(password, "") && ! strcmp(pw_password,""))	~ 1	value of the password	
return 1;		using a salt value	
4: char* encrypted_password = xcrypt(password,	strcmp	Compares two strings and	Yes
pw_password);		returns 0 if the strings	
5: log_user_action(authctxt->user);		match	
6: int authenticated = (strcmp(encrypted_password,	log user action	Records the argument to	No
$pw_password) == 0;$		the system log file (e.g.,	
7: return authenticated;		syslog)	
}			

Figure 55: (a) Example code fragment from SSH program and (b) Functions called from within the code fragment and their roles

**Library Functions:** The *sys\_auth\_passwd* function in Figure 55(a) calls four other functions, namely *shadow\_pw, xcrypt, strcmp and log\_user\_action*. The functionality provided by each of these functions is outlined in Figure 55(b). Of the four functions, the first three are trusted (secure), because they manipulate the critical data and hence, are part of the backward slice. We assume that the source code of the trusted functions is available for analysis.

**Attacks:** To illustrate the IFS technique, we consider two attack scenarios on the untrusted *log\_user\_action* function as follows.

(1) **External attack:** We assume that the *log\_user\_action* function contains a formatstring vulnerability, i.e., it invokes *printf*() using the user-name directly as the first argument without specifying a format-string argument. An external attacker exploits this vulnerability to overwrite any memory location in the program.

(2) **Internal attack:** We assume that the *log\_user\_action* function is supplied by a malicious attacker as part of an external library whose source code is not available for the analysis. The attacker can overwrite any location in memory or registers and change the program's control-flow from within the function. This is so *even if the application contains no memory-corruption vulnerabilities in and of itself.* 

## 7.6 IFS IMPLEMENTATION EXAMPLE

This section illustrates the operation of the IFS algorithm given in Section 7.4. Figure 56 shows the code in Figure 55(a) instrumented with the encoding and decoding functions. Recall that the critical variable chosen is *authenticated*. The backward slice of the *authenticated* variable corresponds to the instructions in the program that can potentially influence the variable's value. Since we consider only intra-procedural slices, we are limited to the instructions in the *sys\_auth\_passwd* function. These are shown in green (or light gray) in Figure 56.

The encoding and decoding operations are represented as functions. In Figure 56, the encoding functions (*encode*) are inserted immediately after the statement that produces a value within the backward slice, while the decoding functions (*decode*) are inserted immediately before the statement within the slice that uses the original value of the variable. Each encoding or decoding function is passed as arguments: (1) the number of the program statement in the slice and (2) the value produced by the program statement.

The functions are marked as *volatile* to prevent the compiler from reordering them during its optimization passes.

The state machines derived for the function in Figure 56 are shown in Figure 57. The accepting states of the state machine are states 2 and 6 (shown in green or as light colored ovals). In the figure, the states in the state machines correspond one-to-one to the encoding calls in the program and are labeled with the encode call arguments.

From Figure 56 and Figure 57, it should be apparent that under normal (attack-free) operation of the program, every value that is encoded has a corresponding decode function before its use in the function (and vice-versa). Further, the state machine reaches an accepting state before the function exits. Therefore, in normal operation, the instrumentation functions do not raise an alarm or perform a false-detection.

We now consider the operation of the program in Figure 56 under attacks. Recall that the goal of the attacker is to overwrite the value of the *authenticated* variable in the program. We first consider generic attacks, where the attacker is unaware that the IFS scheme is deployed, and then targeted attacks where the attacker is aware of the deployment of the IFS technique and actively tries to evade detection by the inserted checks.



# 7.6.1 Generic Attacks

Generic attacks are those in which the attacker is unaware that the IFS scheme is being deployed. Section 7.10 provides a generic proof of the efficacy of the IFS technique for each class of attacks considered in this section.

**External attacks:** An external attacker attempts to exploit the format string vulnerability in the *log\_user\_action* function by crafting an appropriate input to the program. We consider three kinds of external attacks as follows:

(*E1*) Attacker executes system call to overwrite critical data: Assume that the attacker launches a system call by overwriting the return address on the stack with the address of a system call instruction. The attacker also sets up the frame-pointer on the stack such that the system call is executed with the parameters specified by the attacker. The IFS technique by itself does not prevent the attacker from launching the system call, nor does

it prevent the attacker from overwriting the return address (which does not belong to the backward slice of the critical variable). However, once the system call is executed, any attempt by the attacker to overwrite the critical data from within the system call will be detected. For example, if the attacker tries to execute a file read call with the address of the critical variable (*authenticated*) as an argument, the IFS technique detects this as an attack and halts the program.

#### (E2) Attacker overwrites function-pointers/return address to impact the program's

*critical data:* Let us assume that the attacker overwrites the return address on the stack from within the *log\_user\_action* function. The goal of the attacker here is to make the function return directly to line 7, in effect bypassing the initialization of the authenticated variable in line 6. Let us further assume that the *authenticated* variable is assigned to a non-zero, value<sup>31</sup> prior to line 6. This allows the attacker to falsely authenticate herself/himself to the system. The IFS technique detects the attack as follows: the skipping of line 6 results in a control-flow pattern that does not correspond to any valid path within the backward slice of the critical variable *authenticated*. Hence, the state machine in Figure 57 is not in an accepting state when the end of the function is reached (as the corresponding *encode* operation is also skipped<sup>32</sup>). Consequently, any attempt to use the returned value in the called function results in a failure of the *decode* operation and the attack is detected.

<sup>&</sup>lt;sup>31</sup> This is fairly common for local variables in C, which are assigned to arbitrary values prior to their initialization.

<sup>&</sup>lt;sup>32</sup> Section 7.9.2 presents an in-depth analysis of this particular attack, and explains how the attack is detected by the IFS technique.

(E3) Attacker overwrites the critical variable authenticated or any variable in the backward slice- Assume that the attacker chooses a format string such that the value of the pointer variable encrypted\_password is assigned to the address of pw\_password (or vice-versa). This would cause the strcmp function in line 8 to return the value 0 and hence the authenticated variable will assume the value 1, which is the attacker's goal. The IFS technique detects this attack because all variables within the backward slice are in an encoded form prior to the call to the log\_user\_action function. Overwriting the variable results in an error during the decode operation prior to its use.

**Insider Attacks:** In the case of insider attacks, we assume that the function  $log\_user\_action$  is under the control of the attacker. In this case, the function is considered to be a black box in the sense that it can overwrite any set of variables in the program and jump anywhere in the program (recall that its source code may not be available). We consider three cases, depending on the arguments to the  $log\_user\_action$  function.

(11) The log\_user\_action() function modifies the contents of encrypted\_password: This is not allowed according to the semantics of the sys\_auth\_passwd since the log\_user\_action is passed the authctxt->user pointer which can never point to any variable in the backward slice of authenticated (as determined by the compiler's pointer

analysis<sup>33</sup>). The attack is detected by the encoding and decoding operations inserted by the IFS technique.

The next two attacks are not possible for the example in Figure 55(a). Nonetheless, we consider them to illustrate possible insider attacks that may be launched on code that is slightly different from the one in Figure 55(a).

(12) The log\_user\_action function modifies the contents of encrypted\_password, but it is only allowed to modify the contents of the password variable: In addition to passing the user-name to the log\_user\_action function, assume that the sys\_auth\_passwd function also passes a pointer to the user password (not the system password). Since the password variable has already been used in the backward slice before the call to the log\_user\_action function, the programmer may think that the call is harmless. However, it is possible for the log\_user\_action function to maliciously overwrite the contents of the encrypted\_passwd or pw\_passwd strings to make the strings match. This will also be detected by the IFS technique, as only the arguments to the log\_user\_action function, namely, the authctxt->user and password variables, are decoded prior to the function call. Overwrites of any other variable in the backward slice is detected by the IFS technique.

(13) The log\_user\_action function is allowed to modify encrypted\_passwd by virtue of being passed a pointer to the encrypted\_password variable: The attack is not detected

<sup>&</sup>lt;sup>33</sup> It is assumed that the function cannot access global variables unless they are marked as extern, in which case they are treated as function arguments.
by the IFS technique, as the malicious function belongs to the backward slice of the critical variable. Hence, a malicious update is indistinguishable from a legitimate update through the pointer argument of the function. However, a warning can be raised at compile time whenever a pointer to a variable in the backward slice of the critical variable is passed to an untrusted function (i.e., any function whose source-code is not available). This is outside the scope of the current IFS technique.

#### 7.6.2 Targeted Attacks

Targeted attacks are those in which the attacker is aware of the IFS scheme and actively tries to defeat the protection. We assume that the attacker tries to avoid detection as much as possible.

## *Attack 1:* Attacker corrupts a data value produced in the backward slice and calls the corresponding encoding function with the corrupted value.

The attack is detected if the corruption occurs after the value has been encoded. This is because attempting to call an encoding function that has just been called does not correspond to a valid state transition, unless the function repeats in the state machine. Even in the case that the function repeats, the attacker would need to skip the subsequent calls of the encoding functions in the code. In the example in Figure 56, assume that the attacker corrupts the value of  $pw_password$  after the call to the encoding function  $encode_2$  and calls the  $encode_2$  function one more time to re-encode the corrupted value. The attack is detected because the state machine in Figure 57 does not have a transition from state 2 for the  $encode_2$  call. In order to evade detection, the attacker has

to call all the *encode* functions in the state machine following state 2 in Figure 57 or reset the state machine and make it transition through all subsequent states until state 2 is reached. At that point, the call to *encode\_2* becomes valid again and the attacker escapes detection. However, this increases the chances of the attack being detected by other means e.g., timing-based techniques.

# *Attack 2:* Attacker corrupts a data value and manages to bypass the call to the decoding function prior to using the decoded operand within the backward slice of a critical variable

The attack is detected at the next use of the operand in the program by the call to the decoding function. To get away without being detected, the attacker needs to bypass **all** calls to the decoding functions before the variable is used in the backward slice. He/she also needs to ensure that in bypassing the calls to the decoding functions, the encoding functions are not bypassed, as this is detected by the state machine transitioning to an invalid state or not being in an accepting state prior to the function's return.

In the example in Figure 56, assume that the attacker corrupts the value of *authenticated* just before line 7 and bypasses all program statements that subsequently use this operand, including calls to the decode function. The attacker would be able to get away undetected because no checks are performed on the corrupted value. However, in carrying out the attack, the attacker cannot execute any of the code that subsequently uses the value of the *authenticated* variable. This may result in large-scale deviations from the control-flow of

the original program and can be detected by alternate techniques, e.g. control-flow checking [185].

Consider an attack where the attacker changed the value of *pw\_password* after it was encoded and managed to bypass the call to *decode\_2* in line 3. The attack is detected in line 5 when the call to *decode\_2* is encountered again. Bypassing this second call affects the control-flow of the program and results in the call *encode\_6* being skipped in line 6. This attack is detected because the corresponding state machine is not in an accepting state when the function exits i.e., the path is an invalid program path.

# *Attack 3:* The attacker does a replay attack i.e. he/she executes the program, observes a sequence of valid transitions and replaces a run of the state machine with the observed sequence of calls to *encode*.

This attack is detected by a duplication-based encoding scheme (assumed in this chapter), but may bypass other static encoding schemes. In order to detect the attack, the encoding function must be based on a random seed that is chosen at application load time, i.e., each invocation of the program produces a different encoding based on the chosen seed. The randomization ensures that the attacker is not be able to replace a sequence of encoding calls with one from a different instance of the program's execution (unless the seed is the same).

#### 7.7 DISCUSSION

Any static analysis technique, including the IFS, must necessarily approximate the behavior of the program in order to be practically realizable. Typically, the analysis technique over-approximates the behavior of the program, which in turn leads to falsenegatives i.e., missed attacks. This section analyzes the effects of approximations made by the IFS technique on its security guarantees. Unlike the previous section, we do not consider specific attacks mounted by the attacker, but frame the discussion in a more general context.

#### 7.7.1.1 Effect of Intra-procedural Slicing

The IFS technique considers only intra-procedural slices, i.e. it truncates the slice at the beginning of functions. Hence, any corruption of the slice prior to the function call will not be detected by the technique. However, there are two ways of mitigating the impact of intra-procedural slicing as follows:

- Function inlining: This involves inlining the body of the called function into the caller, so that the caller and the callee are treated as one function. This has practical limitations in terms of handling large functions and recursive calls in the code, but does not require programmer intervention beyond specifying the critical variables in the program.
- Choosing critical variables in each function: The user can choose variables in each function such that the entire backward slice is covered. This requires understanding of the dependencies across functions and specifying the critical variables in each function.

The above problems can be solved by considering inter-procedural slices, i.e. contextsensitive dependence analysis. The issue of context-sensitivity is an important one that needs to be addressed by static analysis. However, the issue of context-sensitivity is orthogonal to the IFS technique and is not considered in this chapter.

#### 7.7.1.2 Effect of Acyclic Paths

The IFS technique extracts acyclic control paths in the application and converts the sequences of calls to the encoding functions on the paths into a state machine. Loops in the backward slice of critical variables are represented as cycles in the state machine. However, the state-machines do not include information about the number of iterations executed by a loop. This may be exploited by an attacker who may make the loop execute for fewer or greater number of iterations than allowed by the source program (the attacker's intent may be to bypass security checks performed in loop iterations, or to introduce a semantic violation). Detecting such attacks requires timing/semantic information about program loops or including the loop-counter in the backward slice.

#### 7.7.1.3 Pointer Approximations

When an instruction accesses memory through a pointer variable, the compiler needs to compute the set of locations read/written by the instruction. This set is typically an over-approximation of the set of locations read/written to by the instruction at runtime, and is known as the points-to set of the instruction [101]. The proposed technique relies on the compiler's inferred points-to set for extracting the backward slices. An attacker can replace the memory address used in an instruction (belonging to the backward slice) with another address in the instruction's points-to-set. The attack will be detected only if either the replacement is carried out from an instruction that is not in the backward slice or the

instruction performing the replacement causes the state machine to follow invalid paths or perform illegal transitions.

#### 7.7.1.4 Window of Vulnerability between Encoding Function Calls and Instructions

Any software-based checking scheme will have a window of vulnerability which is the time between carrying out the security check and actually carrying out the privileged operation. If an attack is mounted during this window, it may escape undetected (also known as TOCTTOU vulnerabilities). In the implementation of the IFS scheme, the encoding functions are introduced immediately after the instructions producing the operand, and the decoding functions are introduced immediately before the instructions using the operand. This ensures that the window of vulnerability of the operand is as narrow as possible. In reality, the introduction of encode and decode functions is done at the compiler's intermediate code level, and the code generator may introduce multiple instructions between the encoding/decoding calls and the operands they protect. Furthermore, the compiler may reorder the instructions around the calls to the encoding/decoding functions, leading to dilation of the vulnerability window. Marking the functions as volatile as done by the technique prevents their reordering, but does not alleviate the code-generation problem.

#### 7.7.1.5 Legal but Invalid Control-flow Paths in the Program

Finally, the technique only checks if a sequence of calls to the encoding function is a legal one in the program's control-flow graph. It is possible for an attacker to replace a sequence of encoding calls with a legal but invalid sequence in the program. However, to

avoid detection the sequence must be replaced in its entirety, which may be much harder to achieve for the attacker.

**Summary:** The approximations made by the compiler can have a non-negligible impact on the security guarantees provide by the IFS technique. This is so for any security technique that relies on static analysis, for example the WIT technique [171]. However, the IFS technique differs from these other techniques in two significant ways. First, only the approximations made by the compiler that pertain to the backward slice(s) of the critical variable(s) affect the security guarantees provide by the technique. Secondly, the technique does not need to analyze the code of modules that are not allowed to modify the critical data in the program. This simplifies the code base that must be analyzed statically and hence, the resulting code can be analyzed with higher accuracy.

#### 7.8 EXPERIMENTAL SETUP

**Implementation:** The IFS technique has been implemented as a new pass in the LLVM compiler [99] called the IFS pass. The IFS pass is executed after the lexing, parsing and intermediate representation phases of the compiler, but prior to the register-allocation and code-generation phases. The pass extracts the backward slices of critical data and assigns a unique identifier to each slice instruction. The identifiers of the instructions along different paths in the function are written to a text file. The text file is parsed by Python scripts that generate custom C code to implement the state-machines<sup>34</sup>. The generated C

<sup>&</sup>lt;sup>34</sup> LLVM's intermediate representation is strongly-typed, hence the program types of the instructions are taken into account when generating custom code.

code is linked with runtime libraries for the *encode* and *decode* operations. The IFS pass consists of about 1500 lines of C++ code and the scripts constitute about 500 lines of Python code. The runtime libraries constitute less than 100 lines of uncommented C code.

**Benchmarks:** We demonstrate the IFS technique on server applications. This is because server applications are (1) typically executed with super-user privileges, which makes them extremely attractive targets for attackers, (2) often organized as separate software modules, each of which performs a specific function in the program (for example, the authentication module is responsible for ensuring that only legitimate users are able to gain access to the system) and, (3) consist of different modules executing in a single address space, which allows a malicious module to infiltrate security-critical modules in the application. The server applications considered are as follows:

- OpenSSH: Implementation of the Secure Shell (SSH) protocol. Consists of over 50000 lines of C code [137]
- (2) WuFTP: Implementation of the File Transfer Protocol (FTP), consisting of over 25000 lines of C code [186].
- (3) NullHTTP: a small and efficient multithreaded HTTP server . Consists of about 2500 lines of C code [187].

**Modules:** In the case of the OpenSSH and WuFTP applications, we focus on the security-critical modules of the application, whereas for NullHTTP, we protect the entire application. For OpenSSH, we protect the authentication module while for WuFTP, we protect the user login module (includes authentication and permission checking). In order to facilitate the analysis by the IFS technique, we extract these modules as standalone programs called stubs. Each stub can be executed independently of the main application,

and completely encapsulates the security-critical functionality of the module. The stub for OpenSSH consists of about 250 lines of code, while for WuFTP, the stub consists of about 500 lines of code. The overheads for these two applications are reported in terms of the execution time of the stubs. *Note that the stubs have a much smaller execution time compared to the entire application, and hence better represent the performance overhead of the IFS technique.* (The overheads of the IFS technique were too low to measure in the context of the entire OpenSSH and WuFTP programs). For NullHTTP, we report overheads relative to the application considered as a whole (due to its relatively small size).

**Critical Variables:** In each of the target applications, we choose critical variables based on possible insider attacks that may be launched against the application. The insider attacks considered are as follows:

- **OpenSSH**: The insider allows a colluding user to be authenticated in spite of providing the wrong password.
- **WuFTP:** The insider allows spoofing of a user's identity in order to access the user's files/directories and perform malicious activities so that the user is blamed.
- **NullHTTP:** Attacks that either modify the client request or the response in order to send malicious or unintended content to the user.

Table 32 shows the critical variables in each application that are chosen (manually) in each application.

Application	Critical Variable (Function)	Rationale/Comment
OpenSSH	Return value (auth_password)	Return value is used to decide if user should be authenticated
WuFTP	Return value ( <i>check_Auth</i> )	Return value is used to decide if user should be authenticated
	Resolved_path (wu_real_path)	Stores the home directory of the user to which he/she has access
	user_name (check_Auth)	User name of the user who is attempting to log into the system
NullHTTP	pPostData (doResponse)	Buffer containing client request for processing by the server
	filename (sendFile)	Name of file containing the webpage requested by the client

Table 32: Critical variables in the applications and the rationale for choosing the variables as critical

For the OpenSSH and WuFTP stubs, the LLVM compiler [99] is able to aggressively inline the functions into a single function. Hence, the backward slice of the critical variable encompasses all instructions in the stub. In the NullHTTP application, the LLVM compiler inlines all the functions related to processing a client's request into a single function (*htloop*). The entire backward slices of the critical data are contained entirely within this function. Table 33shows the static characteristics of the inlined function containing the critical variables in each application.

 Table 33: Static characteristics of the instrumentation in each application

Application	Total number of assembly instructions in the function	Number of encode calls in the function	Number of decode calls in the function	Number of acyclic control paths
OpenSSH	2430	252	296	103
WuFTP	800	7	6	2
NullHTTP	5594	404	543	198

**Performance Measurements:** In order to measure the performance overhead of the IFS technique, we executed both the original, non-instrumented program and the instrumented version of the program. The measurements are conducted using the *gettimeofday()* system call on a 2.0 Ghz Pentium 4 Linux system (2 GB RAM).

**Resilience Measurements:** In order to evaluate the resilience of the technique to attacks, we mounted both insider attacks and external attacks on the applications protected with the IFS technique. The insider attacks consisted of replacing a specific function (which does not belong to the backward slice of the critical data) with a malicious surrogate function. The external attacks consisted of planting memory corruption vulnerabilities in the program and exploiting them through specially constructed inputs to impact the critical data in the application.

#### 7.9 EXPERIMENTAL RESULTS

This section presents the results of the experiments evaluating the performance and resilience of the IFS technique. It also examines the reasons for the overheads.

#### 7.9.1 Performance Overheads

**OpenSSH Authentication Module:** In order to evaluate the performance overheads introduced by the IFS technique, the authentication stub is executed with three inputs, consisting of (1) wrong user-name, (2) correct user-name, correct password and, (3) correct user-name, wrong password. In each case, the execution time of the instrumented program is compared to the execution time of the original, non-instrumented version for a given input. The results are shown in Table 34. As observed in the table, the performance overhead ranges from about 47% to 95%, across the inputs, with a mean value of 79 %.

Input	Original (uS)	Instrumented (uS)	Overheads(%)
1	155	228	47
2	159	306	95
3	159	310	95
Mean			79

Table 34: Execution times of SSH authentication stub

**WuFTP Login Module:** We consider three inputs for measuring the performance overhead of the instrumented stub, (1) attempt to log in with the wrong username (username must be ftp) (2) log in with correct username and wrong password, and (3) log in with correct username and correct password. Table 35 shows the execution time overheads for each of the inputs. As seen from the table, the performance overhead of the instrumented application ranges from about 4 % to 11 %, with a mean value of 7.5 %.

Input	Original (uS)	Instrumented (uS)	Overhead(%)
1	45	50	7
2	90	90	4
3	48	53	11
Mean			7.5

Table 35: Execution times of FTP login stub

**NullHTTP Application:** In order to evaluate the performance overhead introduced by the NullHtp server application, we developed a multi-threaded client program (in C) to request web-pages from the server using the HTTP POST command. The NullHTTP server is inherently multi-threaded and hence throughput is a more meaningful measure of performance overhead than latency. The client program spawns multiple threads, each of which sends an HTTP request to the server requesting a given webpage. The NullHTTP server in turn spawns a new thread to handle each incoming connection from the client. We measure the total time at the client to successfully complete execution of all spawned threads. The client executes on the same machine as the server in order to eliminate the effect of network latency in the measurements (as far as possible).

Table 36 summarizes the results for the NullHTTP program. For a single-threaded client, the total time taken for satisfying the request is approximately twice as much for the

instrumented version compared to the original, non-instrumented version. As the number of threads in the client increase, so do the times taken for processing the requests at the server due to the increased workload. However, the overheads steadily decrease from 86 % to 0 % as the number of threads increase from 1 to 25. In the single-threaded case, the performance overhead is dominated by latency, while in the multi-threaded case, the overhead is dominated by the throughput. Thus, the IFS technique has a substantial impact on the latency of the NullHTTP server, but small impact on the throughput.

Table 36: Execution times of NullHTTP program

Threads	Original (uS)	Instrumented (uS)	Overhead (%)
1	3052	5674	86
5	20436	21067	16
10	39057	42324	10
25	100177	101916	0

**Discussion:** From the results above, it can be concluded that the performance overheads are highly dependent on the nature of the application and the choice of critical data. For example, in the SSH stub application, the IFS technique introduces an overhead of nearly 100%, while for FTP, the overhead is less than 10%. The reason for this difference is that in SSH, the backward slice of the critical variable comprises of about 10% of the instructions in the program, while in FTP it comprises less than 1 % of the instructions (see Table 33). Consequently, SSH has a higher number of *encode* and *decode* calls in the program compared to FTP. Further, the number of control-paths that must be tracked by the state machines is much higher in SSH than FTP (100 versus 2). As a result, the state machine for the SSH application has many more states compared to that of the FTP application, and hence incurs higher overhead for tracking state transitions.

The NullHTTP application's characteristics mirror those of the SSH application in terms of the number of instructions in its backward slice (also 10%). Consequently, the performance overheads incurred by the IFS technique for the NullHTTP application (in the single-threaded case) is close to 90 %. However, the overhead can be masked in concurrent requests due to the stateless nature of the HTTP protocol. This in turn allows the threads to be executed in parallel with each other with little or no data sharing among them.

A cursory glance at the results may lead the unsuspecting reader to think that the entire application's execution time is slowed down by a factor of two for the SSH application. However, this is not the case as the above overheads are reported as a fraction of the execution time of the authentication module, which encompasses only a very small fraction of the execution overhead of the application (about 1% or less in a typical user session). Therefore, when evaluated in the context of the entire application, the IFS technique incurs negligible performance overhead for the SSH application. The same reasoning applies for the FTP application. In the case of the NullHTTP application, while a single request may be slowed down by a factor of two, the overall throughput of the server is minimally impacted. Further, the request processing time at the server is only a small fraction of the overall latency experienced by a typical HTTP request which is often routed through multiple network hops.

Analysis of the sources of performance overhead of the instrumentation: In order to understand how to reduce the performance overheads of the IFS technique, it is important to understand the contribution of each instrumentation component added by the IFS

technique. The contribution of a component is measured by commenting out the component from the code generated by the IFS technique, and measuring the execution time with and without the component. The difference in the execution times yields the performance overhead due to the component. The following components are considered in the study:

- **Encode:** Execution of the encoding functions to encode variables within the backward slice
- **Decode:** Execution of the decoding functions to decode encoded variables prior to their use
- **Transitions:** Performing transitions in the state machines depending on the encoding function executed.
- **Checks:** Checking if the state machine is in an accepting state before returns of instrumented functions
- **Memory:** Encoding and decoding of memory objects (immediately after stores and immediately before loads)
- **Other:** Additional instrumentation added by the compiler for support (e.g., stack handling, error reporting)

We consider the OpenSSH application as it incurred the highest performance overheads among the target applications. The performance overhead of each component of the instrumentation is as follows. The calls to the decoding functions have the maximum contribution to the overhead (28%), as these are called 802 times in the program (not to be confused with the static counts in Table 33). This is followed by encoding and decoding of memory operands (21%), as these may necessitate an extra load or store instruction. The state transitions constitute about 19% of the overhead, while the calls to the encode function brings another 15 % (this does not include the overhead of state-machine transitions). This is because the program calls the encoding function 620 times, and each call causes one or more state-machine transitions. The operation to check if a state machine is in an accepting state constitutes only 1% of the overhead, as it is called only at function exits (and there is only one function in the OpenSSH stub after inlining). The category, *other*, constitutes 12 % of the overhead due to the runtime support code added by the IFS pass. An additional 4 % overhead is unaccounted for due to measurement errors.

**Techniques to reduce the performance overhead:** Based on the above results, we see that the encoding and decoding operations constitute the highest overhead among the instrumentation components. The overhead can be alleviated by implementing the encoding and decoding operations in hardware. This will require provision of a high-speed hardware cache to store and retrieve the encoded values on demand. This functionality can also be retrofitted onto existing TPM modules in processors [188]. If implemented as a cache, it will incur nearly zero overheads, and the performance overhead of the technique can be reduced by nearly 65%. Further reduction in the performance overheads (by about 20%) can be achieved by using hardware to implement the state machines for tracking encoding calls in the program. However, this requires the

hardware to be reconfigurable, as each application will have a unique set of state machines that must be configured into the hardware at application load time. The hardware module can be implemented as a part of the Reliability and Security Engine (RSE) [1], which is a hardware framework for executing application-level checks. This is a direction for future investigation.

#### 7.9.2 **Resilience Measurements**

This section discusses the results of experimentally testing the resilience of the SSH stub application (the other applications are not discussed due to lack of space). We first consider the external attacks and then the insider attacks from Section 7.6.1. Recall that in both cases, the attacker's goal is to get the *sys\_auth\_passwd* function to return 1 even if the supplied username and password are not valid.

**External Attack:** As mentioned in Section 7.6.1, we assume that the external attacker exploits the format string vulnerability in the *log\_user\_action* function. The attacker has to supply a well crafted username containing malicious format strings along with an arbitrary password in order to overwrite either the pointer to the *encrypted\_password* (E3) or the return address of *log\_user\_action* (E2). The methodology used to craft the format string is similar in both attacks. Due to the lack of space, only the control-flow attack (E2) is discussed below.

Figure 58(a) shows the stack configuration after *printf* is called within the vulnerable *log\_user\_action* function. The format string corresponding to the attack is shown in Figure 58(b). A typical malicious format string can be divided into three components.

The beginning of the string contains the addresses the attacker wants to overwrite. The second section of the string is generally composed with '%x' parameters in order to increment the internal stack pointer of the format function until it points to the beginning of our format string. The final part of the string contains the paddings and the '%n' parameters that allow us to write desired value<sup>35</sup> into chosen addresses. In order to mount the attack, the attacker needs to determine the following by running the program offline<sup>36</sup>.

- 1. The offset from the bottom of *printf*<sup>2</sup> functions stack frame and the stored format string on the stack.
- 2. The address on the stack of the *log\_user\_action* function's return address.
- 3. The address of the instruction the attacker wants to jump to, i.e. the return statement of *sys\_auth\_passwd*.



Figure 58: (a) Stack layout after the call to printf during the attack and (b) Attacker-supplied format string

<sup>&</sup>lt;sup>35</sup> We are able control the last significant byte of the targeted memory addresses. Therefore, in order to control the value of all four bytes of the targeted memory address the attacker needs to overwrite 4 consecutive addresses shifted by one byte each time.
<sup>36</sup> In order to ensure repeatability of the inputs, we assume that address-space randomization is disabled while carrying out the attacks.

In practice, an attacker may achieve similar results by repeatedly attacking the application with different addresses or through information-leaks in the program.

**Detection:** As explained in Section 7.6.1, the attacker bypasses the *strcmp* function in line 6 (and its corresponding *encode operation*) and jumps directly to the return statement in line 7 (Figure 56). This will cause the state machine to be in a non accepting state when *sys\_auth\_passwd* returns, and the attack will be detected.

Examining the code in Figure 56, it may be thought that the attacker can achieve her goal undetected if she jumps to the *encode(6,authenticated)* statement instead of jumping directly to the return statement. However, this is not the case because the instrumentation is done at the assembly level and an encoding function is inserted after each instruction in the backward slice. The earlier explanation in Section 7.6.1 had coalesced multiple encoding calls at the instruction-level into a single call for simplicity of explanation. Figure 59 shows the relevant assembly code of the instrumented *sys\_auth\_passwd* function to illustrate the above attack. The assembly code corresponds to statement 5 in Figure 56. The instrumentation is such that it is impossible to bypass the *strcmp* function without bypassing the encodings of its arguments *encode(61,pw\_password)* and *encode(62,pw\_password)*.

<pre>1. call <log_user_action>     call <decode(2,pw_password)> 2. push pw_password     call <encode(61,pw_password)>     call <decode(4,encrypted_password)> 3. push encrypted_password     call <encode(62,encrypted_password)>     call <decode(61,pw_password)>     call <decode(61,pw_password)>     call <decode(62,pw_password)>     call <decode(62,pw_password)> 4. call <strcmp> 5. sete %al     call <encode(63 authenticated)=""> </encode(63></strcmp></decode(62,pw_password)></decode(62,pw_password)></decode(61,pw_password)></decode(61,pw_password)></encode(62,encrypted_password)></decode(4,encrypted_password)></encode(61,pw_password)></decode(2,pw_password)></log_user_action></pre>	<pre>log_user_action(char* username[]){     /* regular log_user_action statements*/         //malicious code         if (strcmp(username,"malicious")==0){             ptrEncrytedPw=(char**)username+256;             ptrPwPw=(char**)username+280;             *ptrEncrytedPw=*ptrPwPw;         }     } }</pre>
Figure 59 : Assembly code of the	log_user_action function
instrumented sys_auth_passwd function	

**Insider Attacks:** Section 7.6.1 considers three kinds of insider attacks on the application. Due to the lack of space, we will only discuss attack I1 in detail. For illustration purposes, Figure 60 shows the source code of the malicious *log\_user\_action* function (remember that this code is not available to the IFS technique as *log\_user\_action* is an untrusted function). Since the pointers *encrypted\_password* and *pw\_password* are stored on the stack, it is possible for the malicious library function to corrupt their values by using a fixed offset from *log\_user\_action*'s argument: *username*. Supplying "malicious" as username and an arbitrary (but fixed length) password enables the insider to be authenticated as a legitimate user, thereby achieving the attacker's goal. In practice, this attack is likely to be more subtle as the attacker would try to hide their tracks more cleverly.

**Detection:** The attack is detected by the *decode(4,encrypted\_password)* instruction (Figure 59) in the *sys\_auth\_passwd* function. The value of *encrypted\_password* is checked prior to its use in the *strcmp* function in line 4 (the check is done by the *decode* operation). Any modifications of this value by the *log\_user\_action* will result in a deviation from the variable's value when it was encoded prior to calling the *log\_user\_action* function and the attack will be detected. An analogous argument can be made for corruptions of the *pw\_password variable*.

#### 7.10 PROOF OF EFFICACY OF THE IFS TECHNIQUE

This section provides a semi-formal proof of the efficacy of the IFS technique against both insider attacks as well as external memory corruption attacks. We show that the IFS technique can detect any attempt by an untrusted third-party module to overwrite critical data (independent of whether the module's source code is available) in violation of its expected behavior in the application. We also show that the IFS technique detects memory corruption attacks that result in execution of unwanted system-calls, violations of the application's control flow, or overwriting of security critical data in the program. For each external attack category, we show that the IFS technique detects all attacks that would be detected by state of the art security techniques.

We first discuss insider attacks and then discuss memory corruption attacks launched by an external attacker.

**1.** Attacks launched by an insider: In this case, we assume that an untrusted third-party module is loaded into the same address space as the application. The module may modify the application's data, change its control-flow or execute system calls on behalf of the application. We assume that (1) the module may not modify the application's code, (2) it cannot hook into any system calls made by the application, and (3) the IFS checks themselves cannot be bypassed en-masse (though the individual encode/decode operations may be bypassed).

A function call is legitimately allowed to modify a variable if and only if (1) The variable is declared as a global variable in the program, or (2) The function is passed a pointer that may potentially alias the variable, or (3) The variable is dynamically allocated on the heap and can be reached through a pointer passed to the function. An assignment of the variable to a function's return value is treated as a definition of the variable in the calling function rather than as a modification in the called function.

We consider three cases when an untrusted third-party function is invoked by the application. The cases correspond to whether the source code of the third-party function is available for analysis and whether the function is legitimately allowed to influence the critical data's value according to the C language semantics.

# Case 1: The source code of the third-party function is NOT available AND the function is NOT allowed to legitimately modify any variable in the backward slice of the critical variable.

**Proof:** Since the function is not legitimately allowed to modify variables in the backward slice of critical data, the IFS technique will not insert calls to the encode operation after the function call to re-encode the modified data in the slice. If the function does attempt to modify any of the encoded data, it will violate the encoding, which will be detected when the modified data is decoded prior to its use (within the function). It is also possible for the function to modify the control-flow of its calling function by overwriting its return address. These will be detected by the state-machine transitions of the IFS if the modified control-flow impacts the critical data in any way.

Case 2: The source code of the third party function is NOT available AND the function IS allowed to legitimately modify one or more variables in the backward slice of the critical variable.

**Proof:** The function is legitimately allowed to modify one or more variables in the backward slice (as determined by the static analysis) and hence the variables will be decoded prior to the function call. After the function returns, the variables are re-encoded by the IFS technique. Any modification made by the function to the variables is reflected in the values used in the slice. Note that this is a potential security hazard as the function could perform unknown operations on the variables in the backward slice, thereby influencing the critical variable. However, only modifications to the variables that the function is legitimately allowed to modify are reflected in the program – all other modifications are detected by the IFS technique prior to their use in the program. Further, a compile-time warning is emitted if the backward slice includes a function for which the source code is not available, and the programmer must explicitly override the warning (after presumably vetting that the function is indeed doing the right thing for the values it is allowed to modify in the slice). Also, any modifications of the program's control-flow by the untrusted function are treated similar to violations of control-flow by external attackers and are hence detected by the IFS (provided the modifications impact the critical data either directly or indirectly).

In the above case, since the source code of the function is not available, the function as a whole is assumed to modify the backward slice variables that it is legitimately allowed to modify. In other words, the individual instructions that actually modify the variables are not bracketed with encode and decode operations. Hence, an attacker may be able to modify the data through a different instruction than the one that was not allowed to do so

(within the function). However, the modification needs to be in the set of variables that the function is legitimately allowed to modify.

### Case 3: The source code of the third-party function IS available, in which case it is straightforward to check if the function can legitimately modify one or more variables in the backward slice.

**Proof:** It is assumed that the function is trusted, since its source code is available, and hence it can be analyzed (statically) to determine the set of instructions in the function that are legitimately allowed to modify the data in the backward slice of the critical variable. In this case, there is no ambiguity about the data that the function is allowed to modify (subject to the usual sources of imprecision inherent in static analysis). However, the precision in the determination of the instructions that perform the modifications depends on whether the slicing is intra-procedural or inter-procedural. In the case of intra-procedural slicing, the entire function is assumed to modify the variables in the backward slice (that it is legitimately allowed to modify) and this is similar to case 2. In case the slicing is inter-procedural or if the function can be inlined into the calling function, then it is possible to identify the individual instructions in the function that write to the variables in the backward slice of the critical variable. The instructions can be bracketed by calls to the encode and decode operations to ensure that the window of vulnerability of the data in the slice is minimized. Thus, performing intra-procedural analysis weakens the security guarantees with respect to the instructions within the function, but does NOT weaken the security guarantees with respect to the effect of the function on the rest of the application.

2. Memory corruption attacks: These include attacks that do one or more of the following: (a) launch a system call on behalf of the application to execute another process, (b) change the control-flow of the program by executing the attacker's code as a result of overwriting function pointers/return addresses, or (c) overwrite security-critical data in the application either directly or indirectly. The IFS technique will detect all three cases provided they modify the critical data. We consider each case as follows.

## (a) Attacks that launch system-calls on behalf of the application to overwrite the critical data:

The goal is to show that the IFS scheme will detect an anomalous sequence of system calls that ultimately influence the critical data (i.e. by writing to the data directly or indirectly) provided the attack is also detected by existing system-call based detection techniques[174]. In order to keep this discussion at a generic level, we consider an idealized system-call detection technique, which accepts a system-call sequence if and only if the sequence of system-calls corresponds to a valid path in the program (as determined through static analysis).

**Proof Sketch:** Let the anomalous sequence of system calls be  $(S_1 S_2 \dots S_k)$ . We know that this system call sequence is not admitted by the language of valid system call sequences in the system, say S. i.e. there exists NO valid program path which corresponds to the system call sequence  $S = (S_1 S_2 \dots S_k)$ .

We also know that the sequence of system calls impacts the critical variable<sup>37</sup> either directly or indirectly. Consider a system call  $S_i$  in the sequence that impacts the critical variable (there must be at least one such  $S_i$  in the sequence). There are two possible ways in which  $S_i$  can impact the critical variable (illegitimately).

(1)  $S_i$  can write to the memory address containing a variable in the backward slice of the critical variable.

If the write by the system call  $S_i$  is illegitimate, it will be detected by the IFS as the values in the backward slice are in an encoded form prior to the system call. Any overwriting of the values will violate the encoding and will hence be detected when they are used within the program (during the decode operation).

(2)  $S_i$  can modify the control-flow of the program illegitimately to influence the computation of the critical variable. Let us assume that the runtime sequence of encode operations resulting from these actions is 'I'. We know that the system-call sequence S is not a valid one in the program. We need to show that the sequence I does not correspond to a valid IFS in the program.

The proof proceeds by contradiction – we start by assuming that the signature I is a valid one in the program i.e. there exists a program path along which I belongs to the IFS. This implies there exists a valid program path along which the sequence of instructions transcribed by I also belong to the backward slice of the critical variable (CV). Now,

<sup>&</sup>lt;sup>37</sup> We assume that the application has only one critical variable. The extension to multiple critical variables is straightforward.

consider only the system calls in the backward slice of the CV. By definition, this system call sequence must correspond to a valid program path. Further, the calls themselves are a proper sub-sequence of the system-call based signature of the application. Since they correspond to a valid path, the reason that the system-call sequence S is rejected must be because of a sub-sequence S' that does not involve any legitimate writes to the critical variable (either direct/indirect). Hence, the sequence S' can be removed from the signature S to form a new system-call based signature S'' that impacts the critical variable. However, this corresponds to a valid system-call sequence, and is hence not rejected by the system-call based technique. This is a contradiction of our initial assumption that the system-call based detection technique rejects *all* anomalous system-call sequences that impact the critical variable (recall that we considered idealized system-call sequences).

### (b) Attacks that violate the control-flow integrity of the application to influence critical data:

The goal is to show that the IFS scheme will detect attacks that violate the application's control-flow and influence the critical data, provided the attacks are also detected by control-flow checking schemes. Control-flow checking [185]detects attacks that violate the application's control flow in violation of the program semantics. The violation can occur through the injection of new code by the attacker, or by the attacker overwriting either function-pointers or return addresses on the stack.

**Proof Sketch:** Let the valid set of basic-blocks in the program (function) belong to the set  $B = \{ B_1, B_2, ..., B_n \}$ . The Control-Flow Signature (CFS) is a regular expression consisting of some combination of  $B_i$ s. A control-flow checking technique essentially compares the runtime control-flow to the CFS and checks for validity. A valid CFS is defined as a CFS that corresponds to a valid path in the program (function).

Let the set of basic blocks in which the critical variable<sup>38</sup> (CV) is defined be denoted by the set V. Let BS(V) denote the set of basic-blocks which comprise instructions in the backward slice of the CV. Note that this is different from the IFS as the IFS consists of all instructions in the backward slice of the CV, whereas BS(V) only considers the basic blocks containing the instructions. Let's call this the BIFS.

Our goal is to show that no control-flow sequence that is rejected by the CFS but accepted by the BIFS contains instructions that modify the critical variable either directly or indirectly.

As in the previous case, the proof proceeds by contradiction. Let us assume that there exists a control-flow sequence C that is rejected by CFS but accepted by BIFS and containing at least one instruction I that modifies the value of the critical variable (otherwise, the proof is done). There are two cases for the instruction I as follows:

1. *The instruction I was part of the original program*: In this case, 'I' will be a part of the backward slice of the program, and hence its parent basic block B will be a part of the

<sup>&</sup>lt;sup>38</sup> As before, we assume a single critical variable in the program without loss of generality.

BIFS. Further, the paths considered by the IFS technique in constructing state machines (at compile-time) are derived from the program's control-flow graph (CFG), and are hence a subset of the paths present in the CFS. Therefore any sequence involving the block B that registers as a deviation from the CFS will also be registered as a deviation by the BIFS. This implies that the sequence will also be rejected by the BIFS – which is a contradiction of our initial assumption.

**2.** *The instruction I is introduced by the attacker through a code-injection attack*: In this case, 'I' will not be able to modify the value of the critical variable, or any other variable in the backward slice of the critical variable. This is because all the variables will have been encoded when the instruction I is executed. Moreover, any attempt by instruction I to create a value and call an encoding function will result in a violation of the BIFS (because the basic block containing I will not belong to B, the domain set of BIFS). Finally, if I tries to perform a jump to the middle of an existing control-flow sequence in the BIFS, or to truncate the BIFS, it will be treated analogous to case 1 (i.e. any such jump that is detected by the CFS is also detected by the BIFS). Therefore, either I cannot modify the CV or it will result in a sequence that violates the BIFS, which is a contradiction of our assumption.

Hence, the BIFS detects all control-flow attacks that influence the critical variable and are detected by a control-flow checking technique.

(c) Attacks that overwrite the critical variable either directly or indirectly – The goal is to show that the IFS technique is at least as effective as any other memory-safety

checking technique (e.g. the WIT technique [171]) in detecting memory corruption attacks that attempt to overwrite the critical data. Let us assume that the instruction that performs the overwriting is I and that I belongs to the original program (otherwise, it is equivalent to the attacker mounting a code injection attack and this was covered in case b). Assume that the set of variables in the target set of instruction I is given by T. The WIT technique will detect any attempt by I to write to a variable outside the set T<sup>39</sup>. Since we assume that the WIT technique detects this attack, it must follow that the critical variable V is outside the set T. We consider three cases for the instruction I as follows:

i. *I does not belong to the backward slice of the critical variable:* The attack will be detected as all the variables in the backward slice of I will be encoded and any overwriting of them will result in an incorrect value when decoded. In case 'I' attempts to call an encoding function, it will necessarily violate the sequence of state transitions derived by the IFS technique, and the attack is detected.

ii. *I belongs to the backward slice of the critical variable, but is not valid for the current execution path:* Since the backward slice is conservative by definition, it has to include all instructions that could potentially write to the critical variable, even if they are not valid for the current execution (input). This would be detected by the IFS if (and only if) the execution of the instruction results in an invalid state transition in the state machines i.e. there is at least one other instruction in the path that makes the path invalid. This need

<sup>&</sup>lt;sup>39</sup> In reality, WIT offers a much weaker guarantee, namely that 'I' does not write to objects of a different color than itself. Since merging of colors can occur, I can write to an object outside its target set of the same color. Nonethess, we consider a stronger version of the WIT technique's guarantees.

not always be true – for example, if I is the only instruction along the *then* branch of an *if* statement, when the valid execution consists of the *else* statement. However, such cases would not be detected by the WIT technique, as WIT has to conservatively assume that all paths are valid in the program and hence assign the same color to all instructions that can potentially write to the same set of objects (even if they are on different paths). In this case, the guarantees provided by the IFS technique are *stronger* than that of WIT – provided the attack substitutes instructions corresponding to paths consisting of at least two backward-slice instructions.

iii. *I belongs to the backward slice of the critical variable, and is valid for the current execution path* – This attack cannot be detected by the WIT technique as it does not take into account the order of instructions in the program in determining the validity of a write. However, the IFS technique can provide limited protection against this class of attacks, provided the overwriting instruction is executed in an order that is inconsistent with the backward slice i.e. there exists no program path in which the execution of the instruction would result in an order consistent with the control-flow graph of the program. This is because the execution of the instruction has to either trigger an invalid state transition in the state machine or the values written to by the instruction are dynamically dead at the time of the overwriting (and hence the overwriting is benign). However, if this constraint is not satisfied, the IFS technique cannot detect the attack (and neither can the WIT technique).

Thus we see that in the case of memory corruption attacks on the critical data, the IFS technique can detect all attacks that would also be detected by a memory-safety checking technique such as WIT.

**Summary:** Thus we show semi-formally that the IFS technique detects all cases of insider attacks that attempt to modify critical data independent of whether the untrusted module's source code is available for analysis. In the case of external memory-corruption attacks, the IFS technique detects any attempt to impact the security-critical data provided the attack is also detected by existing, state-of-art, security techniques.

#### 7.11 CONCLUSION

This chapter introduced an approach to protect security critical data from insider attacks. The approach leverages existing static analysis techniques to extract the backward slices of critical variables and to convert the slice into a Information-flow Signature (IFS). The IFS is tracked and checked at runtime using automatically generated code. A deviation of the runtime signature from the derived IFS indicates a security attack. We have deployed the technique on three widely-used open-source applications to protect security-critical data and have shown that the technique detects both insider attacks and external memory corruption attacks with low performance overheads.

Future work will involve (1) implementing the encoding and decoding functions in hardware to reduce the performance overheads of the IFS technique and (2) incorporating context-sensitivity in the slicing algorithm to increase the coverage of the technique.

### CHAPTER 8 CONCLUSIONS AND FUTURE WORK 8.1 CONCLUSIONS

This dissertation has demonstrated a unified approach for providing reliability and security to applications in an automated fashion. The approach presented in this dissertation enhances the compiler and the runtime system to derive application-aware error and attack detectors that continuously monitor the application for errors and attacks. The key characteristic of the approach is that it protects critical data in the application from a wide-range of errors and attacks. Critical data is defined as any data in the application, which if corrupted can cause failures with long downtimes or hard-to-detect security compromises.

The errors considered by the approach include hardware transient errors in memory and computation as well as software defects that cause transient data corruptions. Examples of the latter are soft-errors, errors in the processor's control-logic and variation-induced errors. Examples of the latter are memory corruption errors and race conditions. The attacks considered by the approach include external attacks that exploit memory corruption vulnerabilities in the application as well as internal attacks launched by a trusted insider in the same address space as the application.

The dissertation also presents a unified approach to formally model the effects of lowlevel errors and security attacks on the application. The formal approach exhaustively enumerates the effects of errors (attacks) according to a given fault (threat) model and helps in validating the efficacy of the derived detectors.

Finally, the detectors derived using the approach proposed in this dissertation have been implemented using reconfigurable hardware in the context of the Reliability and Security Engine (RSE) [33], which is a hardware framework for executing application-aware checks. The detectors have been prototyped as part of the Trusted Illiac project at UIUC. In summary, application-aware dependability is a viable approach for building highly reliable and secure applications with lower performance overheads compared to traditional dependability techniques such as duplication or type-safety.

#### 8.2 FUTURE WORK

This section provides a roadmap of future work in the directions explored by this dissertation.

**Compilers and program analysis techniques:** Compilers typically focus on optimizing program performance by removing redundancies in the program's source code. This is because redundancies in the program code can result in wasteful computation and consequently loss of performance. However, redundancies also offer advantages in terms of increasing program resilience to hardware and software errors, provided the redundancies can be turned into appropriate runtime checks.

Currently, even if redundancies are present in an application's source code, they offer little benefit to the application as they are not represented as runtime checks. In the future, compilers could convert redundancies in the program into runtime checks for error detection. Compilers can also introduce redundancies in a controlled manner into the original code or avoid removing certain redundancies that were originally present.

However, a judicious trade-off must be achieved between introducing too much redundancy, which would hurt performance, and not introducing any redundancy, which may impact the application's reliability and security.

The technique proposed in Chapter 4 is one way of introducing redundancies in the computation of critical variables in the form of runtime checks. Further, the check is not just a straightforward duplication of the critical variable's computation, but a selectively optimized version and is hence different from the original version. However, the technique did not attempt to explicitly diversify the representation of the check with respect to the original computation. Diversification can lead to detection of permanent hardware errors and software bugs that would not be detected by duplication. This is an area of future investigation.

**Program verification:** Program verification techniques such as theorem proving and model-checking analyze the program's code in order to prove properties about the program with respect to a formal specification. The formal specifications are provided by the programmer and the verification tool attempts to statically establish whether the program satisfies the specification along each program path. As discussed in Section 4.2.1, these techniques are vulnerable to the *feasible path problem*, which lead to exploration of program paths that will never occur during a concrete execution of the program. Recent work has considered the use of dynamic analysis to drive the verification along concrete program execution paths [189, 190]. The main idea in these systems is to piggyback the symbolic exploration of paths by the verification tool onto

the concrete execution of the program and use the information from the concrete execution to prune infeasible paths.

The SymPLFIED technique described in Chapter 5 is a symbolic technique for formally validating fault-tolerance properties of an application. SymPLFIED also faces the feasible path problem as it statically explores the program's error outcomes. This leads to state-space exploration when model-checking large programs. One way of enhancing the analysis is to use runtime information gathered during the dynamic execution of the program to explore only those paths that are likely to be executed during a concrete execution. However, this is different from the techniques proposed in [189, 190], as these techniques do not need to consider the effects of random errors in the state-space exploration. The main challenge introduced by random errors is that the program may follow paths that do not occur in any dynamic error-free execution. Hence, new methods of integrating dynamic execution profiles with symbolic execution are needed.

**Runtime systems for program monitoring:** Runtime monitoring systems provide a flexible method to observe programs and perform adaptations on the fly depending on changes in the requirement or the environment. The techniques proposed in this dissertation also fall under the broad umbrella of runtime monitoring techniques. However, existing runtime monitoring are geared towards detecting specific kinds of errors [90] and security attacks [191]. The technique proposed in this dissertation on the other hand, can detect a broad class of errors and attacks that impact critical variables in the application. The technique can be integrated into a program monitoring framework such as Monitoring Oriented Programming (MOP) [192] to detect generic runtime errors
and security attacks. This offers the advantage that the checks can be expressed in a formal fashion in order to facilitate reasoning about their detection capabilities. Further, the checks can be adaptively enabled or disabled at runtime depending on the prevalence of errors or attacks in the program's environment as well as the maximum performance overhead that the end-user is willing to incur when executing the application.

**Micro-architecture design:** The technique developed in this dissertation uses a combination of software and reconfigurable FPGA (Field-Programmable Gate Array) hardware to execute the derived detectors. In the future, it is possible that the processor itself directly executes the detectors using specialized functional units with no intervention from the software. Each processor could have dedicated functional units to execute specific detectors. The functional units would be configured by the processor manufacturer to include checks used by a standard set of workloads (benchmarks). This is similar to the approach in [193] for accelerating program operations using hardware.

An interesting challenge in this approach is to specify a set of common detector patterns across a range of applications to be configured into hardware. The processor's front-end can automatically identify the specified patterns in the application's binary and transparently schedule the checks onto dedicated functional units with no involvement from the compiler. This approach will completely move the complexity of runtime adaptation to errors and attacks from the software to the hardware. It will also obviate the distribution of multiple versions of an application with different checks for different environments, and instead allow the hardware to transparently control both the degree and nature of the runtime checks that are executed for an application.

295

## REFERENCES

[1] Nakka, N., Z. Kalbarczyk, R.K. Iyer, and J. Xu. An Architectural Framework for Providing Reliability and Security Support. in International Conference on Dependable Systems and Networks. 2004: IEEE Computer Society.

[2] Ross, J.W. and G. Westerman, Preparing for utility computing: The role of IT architecture and relationship management. IBM Systems Journal, 2004. 43(1): p. 5-19.

[3] Rappa, M.A., The utility business model and the future of computing services. IBM Systems Journal, 2004. 43(1): p. 32-42.

[4] Hayes, B., Cloud computing. 2008.

[5] Armbrust, M., A. Fox, R. Griffith, A.D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, and I. Stoica, Above the clouds: A Berkeley view of cloud computing. University of California, Berkeley, Tech. Rep, 2009.

[6] Gray, J. Why do computers stop and what can be done about it. in Symposium on Reliable Distributed Systems. 1986: IEEE.

[7] Sullivan, M. and R. Chillarege. Software defects and their impact on system availability-a study of field failures in operating systems. in Twenty-First Symposium on Fault-Tolerant Computing. 1991.

[8] Lee, I. and R.K. Iyer, Software dependability in the Tandem GUARDIAN system. IEEE Transactions on Software Engineering, 1995. 21(5): p. 455-467.

[9] Spainhower, L. and W. Bartlett, Commercial Fault Tolerance: A Tale of Two Systems. IEEE Transactions on Dependable and Secure Systems, 2004. 1(1): p. 87-96.

[10] Slegel, T.J., I. Robert M. Averill, M.A. Check, B.C. Giamei, B.W. Krumm, C.A. Krygowski, W.H. Li, J.S. Liptay, J.D. MacDougall, T.J. McPherson, J.A. Navarro, E.M. Schwarz, K. Shum, and C.F. Webb, IBM's S/390 G5 Microprocessor Design. IEEE Micro, 1999. 19(2): p. 12-23.

[11] Saggese, G.P. and A. Vetteth, Microprocessor Sensitivity to Failures: Control vs Execution and Combinational vs Sequential Logic, in Proceedings of the 2005 International Conference on Dependable Systems and Networks. 2005, IEEE Computer Society.

[12] Nakka, N., K. Pattabiraman, and R. Iyer. Processor-level Selective Replication. in International Conference on Dependable Systems and Networks (DSN). 2007: IEEE.

[13] Rinard, M., C. Cadar, D. Dumitran, D.M. Roy, T. Leu, and W.S.B. Jr., Enhancing server availability and security through failure-oblivious computing, in Proceedings of the 6th conference on Symposium on Opearting Systems Design \& Implementation - Volume 6. 2004, USENIX Association: San Francisco, CA.

[14] Gu, W., Z. Kalbarczyk, R. Iyer, and Z. Yang. Characterization of linux kernel behavior under errors. in International Conference on Dependable Systems and Networks. 2003: IEEE Computer Society.

[15] Chillarege, R. and R.K. Iyer, Measurement-based analysis of error latency. IEEE Transactions on Computers, 1987. 36(5): p. 529-537.

[16] Goswami, K.K., R.K. Iyer, and L. Young, DEPEND: A Simulation-Based Environment for System Level Dependability Analysis. IEEE Transactions on Computers, 1997. 46(1): p. 60-74.

[17] Pattabiraman, K., Z. Kalbarczyk, and R.K. Iyer. Application-based metrics for strategic placement of detectors. in Pacific Rim Dependable Computing. 2005. Changsha, China: IEEE CS Press.

[18] Weiser, M. Program Slicing. in Fifth International Conference on Software Engineering. 1981: IEEE Computer Society Press.

[19] Hsueh, M.-C., T.K. Tsai, and R.K. Iyer, Fault Injection Techniques and Tools. Computer, 1997. 30(4): p. 75-82. [20] Cowan, C., C. Pu, D. Maier, J. Walpole, P. Bakke, A.G. SteveBeattie, P. Wagle, Q. Zhang, and H. Hinton. StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks. in Seventh Usenix Technical Symposium. 1998: Usenix.

[21] Forrest, S., S.A. Hofmeyr, A. Somayaji, and T.A. Longstaff, A Sense of Self for Unix Processes, in Proceedings of the 1996 IEEE Symposium on Security and Privacy. 1996, IEEE Computer Society.

[22] Ruwase, O. and M.S. Lam. A practical dynamic buffer overflow detector. in 11th Annual Network and Distributed System Security. 2004.

[23] Dhurjati, D., S. Kowshik, and V. Adve. SAFECode: enforcing alias analysis for weakly typed languages. in ACM SIGPLAN conference on Programming language design and implementation. 2006. Ottawa, Ontario, Canada: ACM Press.

[24] Jones, R.W.M. and P.H.J. Kelly. Backwards-Compatible Bounds Checking for Arrays and Pointers in C Programs. in Automated and Algorithmic Debugging. 1997.

[25] Suh, G.E., J.W. Lee, D. Zhang, and S. Devadas. Secure program execution via dynamic information flow tracking. in 11th international conference on Architectural support for programming languages and operating systems. 2004. Boston, MA, USA: ACM Press.

[26] Chen, S., N. Nakka, J. Xu, Z. Kalbarczyk, and R. Iyer. Defeating Memory Corruption Attacks via Pointer Taintedness Detection. in International Conference on Dependable Systems and Networks. 2005: IEEE Computer Society.

[27] Dalton, M., H. Kannan, and C. Kozyrakis, Raksha: a flexible information flow architecture for software security, in Proceedings of the 34th annual international symposium on Computer architecture. 2007, ACM: San Diego, California, USA.

[28] Xu, J., Z. Kalbarczyk., and R. Iyer. Transparent runtime randomization for security. in 22nd International Symposium on Reliable Distributed Systems. 2003.

[29] Bhatkar, S., D.C. DuVarney, and R. Sekar, Address obfuscation: an efficient approach to combat a board range of memory error exploits, in Proceedings of the 12th conference on USENIX Security Symposium - Volume 12. 2003, USENIX Association: Washington, DC.

[30] Berger, E.D. and B.G. Zorn. DieHard: probabilistic memory safety for unsafe languages. in ACM SIGPLAN conference on Programming language design and implementation. 2006. Ottawa, Ontario, Canada: ACM Press.

[31] Shacham, H., M. Page, B. Pfaff, E.-J. Goh, N. Modadugu, and D. Boneh, On the effectiveness of addressspace randomization, in Proceedings of the 11th ACM conference on Computer and communications security. 2004, ACM: Washington DC, USA.

[32] Sezer, E.C., P. Ning, C. Kil, and J. Xu, Memsherlock: an automated debugger for unknown memory corruption vulnerabilities, in Proceedings of the 14th ACM conference on Computer and communications security. 2007, ACM: Alexandria, Virginia, USA.

[33] Nakka, N., Reliability and Security Engine: A processor-level framework for application-aware detection, in Electrical and Computer Engineering. 2006, UIUC: Urbana-Champaign.

[34] M.Clavel, F.Duran, S.Eker, P.Lincoln, N.Marti-Oliet, J.meseguer, and J.Quesada, Maude: Specification and Programming in Rewriting Logic, in Maude System Documentation. 1999, SRI.

[35] LEON3 Implementation of the Sparc V8. [cited; Available from: <u>http://www.gaisler.com</u>.

[36] Iyer, R.K., D.J. Rossetti, and M.C. Hsueh, Measurement and modeling of computer reliability as affected by system activity. ACM Trans. Comput. Syst., 1986. 4(3): p. 214-237.

[37] Basile, C., W. Long, Z. Kalbarczyk, and R. Iyer. Group communication protocols under errors. in 22nd International Symposium on Reliable Distributed Systems. 2003.

[38] Chandra, S. and P.M. Chen. How Fail-Stop are Faulty Programs? in Proceedings of the The Twenty-Eighth Annual International Symposium on Fault-Tolerant Computing. 1998: IEEE Computer Society.

[39] Bartlett, J.F., A NonStop kernel, in Proceedings of the eighth ACM symposium on Operating systems principles. 1981, ACM: Pacific Grove, California, United States.

[40] Hiller, M., A. Jhumka, and N. Suri. On the Placement of Software Mechanisms for Detection of Data Errors. in International Conference on Dependable Systems and Networks. 2002: IEEE Computer Society.

[41] Jeffrey, M.V. and W.M. Keith, The Avalanche Paradigm: An Experimental Software Programming Technique for Improving Fault-tolerance, in Proceedings of the IEEE Symposium and Workshop on Engineering of Computer Based Systems. 1996, IEEE Computer Society.

[42] Goradia, T., Dynamic impact analysis: a cost-effective technique to enforce error-propagation. SIGSOFT Softw. Eng. Notes, 1993. 18(3): p. 171-181.

[43] Ernst, M.D., J. Cockrell, W.G. Griswold, and D. Notkin. Dynamically discovering likely program invariants to support program evolution. in 21st international conference on Software engineering. 1999. Los Angeles, California, United States: IEEE Computer Society Press.

[44] Narayanan, S.H.K., S.W. Son, M. Kandemir, and F. Li, Using loop invariants to fight soft errors in data caches, in Proceedings of the 2005 conference on Asia South Pacific design automation. 2005, ACM: Shanghai, China.

[45] Benso, A., S. Chiusano, P. Prinetto, and L. Tagliaferri. A C/C++ Source-to-Source Compiler for Dependable Applications. in International Conference on Dependable Systems and Networks (formerly FTCS-30 and DCCA-8). 2000: IEEE Computer Society.

[46] Nethercote, N. and A. Mycroft, Redux: A Dynamic Dataflow Tracer. Electronic Notes on Theoretical Computer Science, 2003. 89(2).

[47] Zhang, X., R. Gupta, and Y. Zhang, Cost and precision tradeoffs of dynamic data slicing algorithms. ACM Trans. Program. Lang. Syst., 2005. 27(4): p. 631-661.

[48] Chillarege, R., W.-L. Kao, and R.G. Condit, Defect type and its impact on the growth curve, in Proceedings of the 13th international conference on Software engineering. 1991, IEEE Computer Society Press: Austin, Texas, United States.

[49] Kao, W.-L, R.K. Iyer, and D. Tang, FINE: A Fault Injection and Monitoring Environment for Tracing the UNIX System Behavior Under Faults. IEEE Trans. Softw. Eng., 1993. 19(11): p. 1105-1118.

[50] Austin, T., E. Larson, and D. Ernst, SimpleScalar: An Infrastructure for Computer System Modeling. Computer, 2002. 35(2): p. 59-67.

[51] Hutchins, M., H. Foster, T. Goradia, and T. Ostrand. Experiments of the effectiveness of dataflow- and controlflow-based test adequacy criteria. in 16th international conference on Software engineering. 1994. Sorrento, Italy: IEEE Computer Society Press.

[52] Bush, W.R., J.D. Pincus, and D.J. Sielaff, A static analyzer for finding dynamic programming errors. Software Practice and Experience, 2000. 30(7): p. 775-802.

[53] Evans, D., J. Guttag, J. Horning, and Y.-M. Tan. LCLint: a tool for using specifications to check code. in 2nd ACM SIGSOFT symposium on Foundations of software engineering. 1994. New Orleans, Louisiana, United States: ACM Press.

[54] Andrews, D. Using executable assertions for testing and fault tolerance, in 9th Faul-tolerance Computing Symposium. 1979. Madison, WI: IEEE.

[55] Leveson, N.G., S.S. Cha, J.C. Knight, and T.J. Shimeall, The use of self checks and voting in software error detection: an empirical study. IEEE Transactions on Software Engineering, 1990. 16(4): p. 432-443.

[56] Hiller, M. Executable Assertions for Detecting Data Errors in Embedded Control Systems. in International Conference on Dependable Systems and Networks (formerly FTCS-30 and DCCA-8). 2000: IEEE Computer Society.

[57] Voas, J., Software testability measurement for intelligent assertion placement. Software Quality Control, 1997. 6(4): p. 327-336.

[58] Patterson, D.A. and J.L. Hennessy, Computer architecture: a quantitative approach. 1990: Morgan Kaufmann Publishers Inc. 594.

[59] Wang, N.J. and S.J. Patel, ReStore: Symptom-Based Soft Error Detection in Microprocessors. IEEE Trans. Dependable Secur. Comput., 2006. 3(3): p. 188-201.

[60] Necula, G.C., S. McPeak, and W. Weimer. CCured: type-safe retrofitting of legacy code. in ACM SIGPLAN-SIGACT symposium on Principles of programming languages. 2002. Portland, Oregon: ACM Press.

[61] Engler, D., D.Y. Chen, S. Hallem, A. Chou, and B. Chelf. Bugs as deviant behavior: a general approach to inferring errors in systems code. in Eighteenth ACM Symposium on Operating systems principles. 2001. Banff, Alberta, Canada: ACM Press.

[62] Hangal, S. and M.S. Lam. Tracking down software bugs using automatic anomaly detection. in 24th International Conference on Software Engineering. 2002. Orlando, Florida: ACM Press.

[63] Maxion, R.A. and K.M.C. Tan, Anomaly Detection in Embedded Systems. IEEE Trans. Comput., 2002. 51(2): p. 108-120.

[64] Rela, M.Z., H. Madeira, and J.G. Silva. Experimental evaluation of the fail-silent behaviour in programs with consistency checks. in Annual Symposium on Fault-tolerant Computing, 1996. Sendai.

[65] Racunas, P., K. Constantinides, S. Manne, and S.S. Mukherjee, Perturbation-based Fault Screening, in Proceedings of the 2007 IEEE 13th International Symposium on High Performance Computer Architecture. 2007, IEEE Computer Society.

[66] Dimitrov, M. and H. Zhou, Unified Architectural Support for Soft-Error Protection or Software Bug Detection, in Proceedings of the 16th International Conference on Parallel Architecture and Compilation Techniques. 2007, IEEE Computer Society.

[67] Sahoo, S., M.-l. Li, P. Ramachandran, V.S. Adve, S.V. Adve, and Y. Zhou. Using likely program invariants to detect hardware errors. in Proceedings of International Conference on Dependable Systems and Networks (DSN). 2008. Anchorage, AK: IEEE.

[68] Oh, N., P.P. Shirvani, and E.J. McCluskey, Error detection by duplicated instructions in super-scalar processors. IEEE Transactions on Relibility, 2002. 51(1): p. 63-75.

[69] Reis, G.A., J. Chang, N. Vachharajani, R. Rangan, and D.I. August. SWIFT: Software Implemented Fault Tolerance. in International symposium on Code generation and optimization. 2005: IEEE Computer Society.

[70] Iyer, R.K., N.M. Nakka, Z.T. Kalbarczyk, and S. Mitra, Recent advances and new avenues in hardware-level reliability support. Micro, IEEE, 2005. 25(6): p. 18-29.

[71] Das, M., S. Lerner, and M. Seigle. ESP: path-sensitive program verification in polynomial time. in Proceedings of the ACM SIGPLAN 2002 Conference on Programming language design and implementation. 2002. Berlin, Germany: ACM Press.

[72] Ball, T. and S. Rajamani. The SLAM Toolkit. in 13th International Conference on Computer Aided Verification. 2001: Springer-Verlag.

[73] Flanagan, C. and S. Qadeer. Predicate abstraction for software verification. in 29th ACM SIGPLAN-SIGACT symposium on Principles of programming languages. 2002. Portland, Oregon: ACM Press.

[74] Demsky, B., M.D. Ernst, P.J. Guo, S. McCamant, J.H. Perkins, and M. Rinard. Inference and enforcement of data structure consistency specifications. in International symposium on Software testing and analysis. 2006. Portland, Maine, USA: ACM Press.

[75] Li, Z. and Y. Zhou. PR-Miner: automatically extracting implicit programming rules and detecting violations in large software code. in 13th ACM SIGSOFT international symposium on Foundations of software engineering. 2005. Lisbon, Portugal: ACM Press.

[76] Mukherjee, S.S., M. Kontz, and S.K. Reinhardt. Detailed design and evaluation of redundant multithreading alternatives. in 29th annual international symposium on Computer architecture. 2002. Anchorage, Alaska: IEEE Computer Society.

[77] Sundaramoorthy, K., Z. Purser, and E. Rotenburg. Slipstream processors: improving both performance and fault tolerance. in Ninth international conference on Architectural support for programming languages and operating systems. 2000. Cambridge, Massachusetts, United States: ACM Press.

[78] Avizenies, A., The Methodology of N-Version Programming, in Software Fault Tolerance, M.R. Lyu, Editor. 1995, Wiley. p. 23-46.

[79] Knight, J.C. and N.G. Leveson, An experimental evaluation of the assumption of independence in multiversion programming. IEEE Transactions on Software Engineering, 1986. 12(1): p. 96-109.

[80] Ammann, P.E. and J.C. Knight, Data Diversity: An Approach to Software Fault Tolerance. IEEE Transactons on Computers, 1988. 37(4): p. 418-425.

[81] Oh, N., S. Mitra, and E.J. McCluskey, ED<sup>4</sup>I: error detection by diverse data and duplicated instructions. IEEE Transactions on Computers, 2002. 51(2): p. 180-199.

[82] Jonathan, C., A.R. George, and I.A. David. Automatic Instruction-Level Software-Only Recovery. in International Conference on Dependable Systems and Networks (DSN'06). 2006: IEEE Computer Society.

[83] Proudler, I.K., Idempotent AN codes, in IEE Colloquium on Signal Processing Applications of Finite Field Mathematics. 1989. p. 8/1-8/5.

[84] Dhurjati, D. and V. Adve. Backwards-compatible array bounds checking for C with very low overhead. in 28th international conference on Software engineering. 2006. Shanghai, China: ACM Press.

[85] Savage, S., M. Burrows, G. Nelson, P. Sobalvarro, and T. Anderson, Eraser: a dynamic data race detector for multithreaded programs. ACM Transactions on Computer Systems, 1997. 15(4): p. 391-411.

[86] Alkhalifa, Z., V.S.S. Nair, N. Krishnamurthy, and J.A. Abraham, Design and Evaluation of System-Level Checks for On-Line Control Flow Error Detection. IEEE Transactions on Parallel and Distributed Systems, 1999. 10(6): p. 627-641.

[87] Bagchi, S., Y. Liu, K. Whisnant, Z. Kalbarczyk, R. Iyer, Y. Levendel, and L. Votta. A framework for database audit and control flow checking for a wireless telephone network controller. in Dependable Systems and Networks. 2001. Goteborg: IEEE CS Press.

[88] Oh, N., P.P. Shirvani, and E.J. McCluskey, Control-flow checking by software signatures. IEEE Transactions on Reliability, 2002. 51(1): p. 111-122.

[89] Havelund, K. and G. Rosu, An Overview of the Runtime Verification Tool Java PathExplorer. Formal Methods in System Design, 2004. 24(2): p. 189-215.

[90] Kim, M., M. Viswanathan, S. Kannan, I. Lee, and O. Sokolsky, Java-MaC: A Run-Time Assurance Approach for Java Programs. Formal Methods in System Design, 2004. 24(2): p. 129-155.

[91] Lattner, C. and V. Adve. Automatic pool allocation: improving performance by controlling data structure layout in the heap. in ACM SIGPLAN conference on Programming language design and implementation. 2005. Chicago, IL, USA: ACM Press.

[92] Ohlsson, J., M. Rimen, and U. Gunneflo. A study of the effects of transient fault injection into a 32-bit RISC with built-in watchdog. in Twenty-Second International Symposium on Fault-tolerant Computing. 1992: IEEE.

[93] Saib, S.H. Executable Assertions - An Aid To Reliable Software. in 11th Asilomar Conference Circuits Systems and Computers. 1978.

[94] Kuang-Hua, H. and J.A. Abraham, Algorithm-Based Fault Tolerance for Matrix Operations. IEEE Transactions on Computers, 1984. C-33(6): p. 518-528.

[95] Avizienis, A., J.C. Laprie, B. Randell, and C. Landwehr, Basic concepts and taxonomy of dependable and secure computing. IEEE Transactions on Dependable and Secure Computing, 2004. 1(1): p. 11-33.

[96] Tip, F., A survey of program slicing techniques. Journal of Programming Languages, 1995. 3(3): p. 121-189.

[97] Kernighan, B.W. and D.M. Ritchie, The C Programming Language. 1989: Prentice Hall Press.

[98] Engler, D. and K. Ashcraft, RacerX: effective, static detection of race conditions and deadlocks. SIGOPS Oper. Syst. Rev., 2003. 37(5): p. 237-252.

[99] Lattner, C. and V. Adve. LLVM: A Compilation Framework for Lifelong Program Analysis \& Transformation. in international symposium on Code generation and optimization. 2004. Palo Alto, California: IEEE Computer Society. [100] Cytron, R., J. Ferrante, B.K. Rosen, M.N. Wegman, and F.K. Zadeck, Efficiently computing static single assignment form and the control dependence graph. ACM Transactions on Programming Languages and Systems, 1991. 13(4): p. 451-490.

[101] Muchnick, S.S., Advanced compiler design and implementation. 1997: Morgan Kaufmann Publishers Inc. 856.

[102] Weicker, R.P., An Overview of Common Benchmarks, in Computer. 1990. p. 65-75.

[103] Carlisle, M.C. and A. Rogers. Software caching and computation migration in Olden. in Fifth ACM SIGPLAN symposium on Principles and practice of parallel programming. 1995. Santa Barbara, California, United States: ACM Press.

[104] Meixner, A. and D.J. Sorin, Error Detection Using Dynamic Dataflow Verification, in Proceedings of the 16th International Conference on Parallel Architecture and Compilation Techniques. 2007, IEEE Computer Society.

[105] Meixner, A., M.E. Bauer, and D.J. Sorin, Argus: Low-Cost, Comprehensive Error Detection in Simple Cores. IEEE Micro, 2008. 28(1): p. 52-59.

[106] Ball, T. and J.R. Larus, Efficient path profiling, in Proceedings of the 29th annual ACM/IEEE international symposium on Microarchitecture. 1996, IEEE Computer Society: Paris, France.

[107] Vaswani, K., M.J. Thazhuthaveetil, and Y.N. Srikant, A Programmable Hardware Path Profiler, in Proceedings of the international symposium on Code generation and optimization. 2005, IEEE Computer Society.

[108] Zhang, T., X. Zhuang, S. Pande, and W. Lee, Anomalous path detection with hardware support, in Proceedings of the 2005 international conference on Compilers, architectures and synthesis for embedded systems. 2005, ACM: San Francisco, California, USA.

[109] Pattabiraman, K., G.P. Saggese, D. Chen, Z. Kalbarczyk, and R.K. Iyer. Dynamic Derivation of Application-Specific Error Detectors and their Implementation in Hardware. in Sixth European Dependable Computing Conference. 2006. Coimbra, Portugal: IEEE CS Press.

[110] Pattabiraman, K. and R. Iyer, Automated Derivation of Application-Aware Error Detectors using Static Analysis. 2006, University of Illinois (Urbana-Champaign).

[111] Perry, F., L. Mackey, G.A. Reis, J. Ligatti, D.I. August, and D. Walker, Fault-tolerant typed assembly language, in Proceedings of the 2007 ACM SIGPLAN conference on Programming language design and implementation. 2007, ACM: San Diego, California, USA.

[112] Clarke, E.M. and B.-H. Schlingloff, Model checking, in Handbook of automated reasoning. 2001, Elsevier Science Publishers B. V. p. 1635-1790.

[113] Larsson, D. and R. Hahnle, Symbolic Fault Injection. International Verification Workshop (VERIFY), International Conference on Automated Deduction (CADE), 2007. 259: p. 85-103.

[114] Arora, A. and S.S. Kulkarni, Detectors and Correctors: A Theory of Fault-Tolerance Components, in Proceedings of the The 18th International Conference on Distributed Computing Systems. 1998, IEEE Computer Society.

[115] Jhumka, A., M. Hiller, V. Claesson, and N. Suri, On systematic design of globally consistent executable assertions in embedded software, in Proceedings of the joint conference on Languages, compilers and tools for embedded systems: software and compilers for embedded systems. 2002, ACM: Berlin, Germany.

[116] Nicolescu, B., N. Gorse, Y. Savaria, E.M. Aboulhamid, and R. Velazco, On the use of model checking for the verification of a dynamic signature monitoring approach. IEEE Transactions on Nuclear Science, 2005. 52(5(2)): p. 1555-1561.

[117] King, J.C., Symbolic execution and program testing. Commun. ACM, 1976. 19(7): p. 385-394.

[118] Boyer, R.S. and J.S. Moore, Program verification. J. Autom. Reason., 1985. 1(1): p. 17-23.

[119] Chen, H., D. Dean, and D. Wagner, Model-checking one million lines of C code. Network and Distributed System Security Symposium, 2004: p. 171-185.

[120] Srivas, M. and M. Bickford, Formal Verification of a Pipelined Microprocessor. IEEE Softw., 1990. 7(5): p. 52-64.

[121] Krautz, U., M. Pflanz, C. Jacobi, H.W. Tast, K. Weber, and H.T. Vierhaus, Evaluating coverage of error detection logic for soft errors using formal methods, in Proceedings of the conference on Design, automation and test in Europe: Proceedings. 2006, European Design and Automation Association: Munich, Germany.

[122] Seshia, S.A., W. Li, and S. Mitra, Verification-guided soft error resilience, in Proceedings of the conference on Design, automation and test in Europe. 2007, EDA Consortium: Nice, France.

[123] Clavel, M., F. Duran, S. Eker, P. Lincoln, N. Marti-Oliet, J. Meseguer, and C. Talcott. The Maude 2.0 System. in Rewriting Technologies and Applications. 2001: Springer.

[124] Serbanuta, T.F., G. Rosu, and J. Meseguer, A Rewriting Logic Approach to Operational Semantics (Extended Abstract). Electron. Notes Theor. Comput. Sci., 2007. 192(1): p. 125-141.

[125] Administration, F.A., TCAS II Collision Avoidance System (CAS) System Requirements Specification.1993.

[126] Lygeros, J. and N. Lynch, On the formal verification of the TCAS conflict resolution algorithms. In Proceedings of the 36th IEEE Conference on Decision and Control, 1997: p. 1829--1834.

[127] Coen-Porisini, A., G. Denaro, C. Ghezzi, and M. Pezz, Using symbolic execution for verifying safety-critical systems. SIGSOFT Softw. Eng. Notes, 2001. 26(5): p. 142-151.

[128] Anderson, R.J., P. Beame, S. Burns, W. Chan, F. Modugno, D. Notkin, and J.D. Reese, Model checking large software specifications. SIGSOFT Softw. Eng. Notes, 1996. 21(6): p. 156-166.

[129] Randazzo, M.R., M.M. Keeney, E.F. Kowalski, D.M. Cappelli, and A.P. Moore, Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector. 2004, U.S. Secret Service and CERT Coordination Center/Software Engineering Institute: Philadelphia, PA. p. 25.

[130] Keeney, M.M. and E.F. Kowalski, Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors. 2005, CERT/CC: Philadelphia, PA.

[131] Polk, W. and L. Bassham, Threat assessment of malicious code and human threats. 1994, NIST, Computer Security Division.

[132] Sheyner, O., J. Haines, S. Jha, R. Lippmann, and J.M. Wing, Automated Generation and Analysis of Attack Graphs, in Proceedings of the 2002 IEEE Symposium on Security and Privacy. 2002, IEEE Computer Society.

[133] Ammann, P., D. Wijesekera, and S. Kaushik, Scalable, graph-based network vulnerability analysis, in Proceedings of the 9th ACM conference on Computer and communications security. 2002, ACM: Washington, DC, USA.

[134] Phillips, C. and L.P. Swiler, A graph-based system for network-vulnerability analysis, in Proceedings of the 1998 workshop on New security paradigms. 1998, ACM: Charlottesville, Virginia, United States.

[135] Probst, C.W., R.R. Hansen, and F. Nielson, Where can an Insider Attack ?, in Formal Aspects in Security and Trust. 2007, Springer Berlin / Heidelberg. p. 127-142.

[136] Pattabiraman, K., N. Nakka, and Z. Kalbarczyk. SymPLFIED: Symbolic Program Level Fault-Injection and Error-Detection Framework. in International Conference on Dependable Systems and Networks (DSN). 2008.

[137] OpenSSH Development Team., OpenSSH 4.21. 2004.

[138] King, S.T. and P.M. Chen, Backtracking intrusions. SIGOPS Oper. Syst. Rev., 2003. 37(5): p. 223-236.

[139] Musuvathi, M. and D.R. Engler, Model checking large network protocol implementations, in Proceedings of the 1st conference on Symposium on Networked Systems Design and Implementation - Volume 1. 2004, USENIX Association: San Francisco, California.

[140] Chinchani, R., A. Iyer, H.Q. Ngo, and S. Upadhyaya, Towards a Theory of Insider Threat Assessment, in Proceedings of the 2005 International Conference on Dependable Systems and Networks. 2005, IEEE Computer Society.

[141] Costa, M., M. Castro, L. Zhou, L. Zhang, and M. Peinado, Bouncer: securing software by blocking bad input, in Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles. 2007, ACM: Stevenson, Washington, USA.

[142] Molnar, D.A. and D. Wagner, Catchconv: Symbolic execution and run-time type inference for integer conversion errors. 2007, EECS Department, University of California, Berkeley: Berkeley, CA.

[143] Kruegel, C., E. Kirda, D. Mutz, W. Robertson, and G. Vigna, Automating mimicry attacks using static binary analysis, in Proceedings of the 14th conference on USENIX Security Symposium - Volume 14. 2005, USENIX Association: Baltimore, MD.

[144] Cadar, C., V. Ganesh, P.M. Pawlowski, D.L. Dill, and D.R. Engler, EXE: automatically generating inputs of death, in Proceedings of the 13th ACM conference on Computer and communications security. 2006, ACM: Alexandria, Virginia, USA.

[145] Boneh, D., R. DeMillo, and R.J. Lipton. On the importance of checking crypto-graphic protocols for faults. in Advances in Cryptgraphy (EuroCrypt). 1997: Springer.

[146] Voyiatzis, A.G. and D.N. Serpanos, Active Hardware Attacks and Proactive Countermeasures, in Proceedings of the Seventh International Symposium on Computers and Communications (ISCC'02). 2002, IEEE Computer Society.

[147] Xu, J., S. Chen, Z. Kalbarczyk, and R.K. Iyer, An Experimental Study of Security Vulnerabilities Caused by Errors, in Proceedings of the 2001 International Conference on Dependable Systems and Networks (formerly: FTCS). 2001, IEEE Computer Society.

[148] Govindavajhala, S. and A.W. Appel, Using Memory Errors to Attack a Virtual Machine, in Proceedings of the 2003 IEEE Symposium on Security and Privacy. 2003, IEEE Computer Society.

[149] Hissam, S., D. Plakosh, and C. Weinstock. Trust and Vulnerability in Open Source Software. in Proceedings - Software Engineering. 2002: IEE.

[150] Jeremey. Linux Kernel Backdoor attempt. 2003 [cited; Available from: http://kerneltrap.org/node/1584.

[151] Hunt, F. and P. Johnson. On the Pareto-distribution of Sourceforge projects. in Open Source Software Development Workshop. 2002.

[152] Veracode, Protecting Your Applications from Backdoors: How Static Binary Analysis Helps Build High-Assurance Applications, in White Paper. 2008.

[153] Cappelli, D.M., T. Caron, R.F. Trzeciak, and A.P. Moore, Programming Techniques Used as an Insider Attack Tool, in Spotlight On, CERT, Editor. 2008, Software Engineering Institute (SEI): Pittsburgh, PA.

[154] Skoudis, E. and L. Zeltser, Malware: Fighting Malicious Code. 2003, Upper Saddle River, NJ, USA: Prentice Hall PTR.

[155] Christodorescu, M. and S. Jha, Testing malware detectors. SIGSOFT Softw. Eng. Notes, 2004. 29(4): p. 34-44.

[156] Christodorescu, M., S. Jha, S.A. Seshia, D. Song, and R.E. Bryant. Semantics-aware malware detection. 2005.

[157] Bayer, U., C. Kruegel, and E. Kirda. TTAnalyze: A tool for analyzing malware. 2006.

[158] Dinaburg, A., P. Royal, M. Sharif, and W. Lee. Ether: malware analysis via hardware virtualization extensions. 2008: ACM New York, NY, USA.

[159] Iyer, R.K., Z. Kalbarczyk, K. Pattabiraman, W. Healey, W.-M.W. Hwu, P. Klemperer, and R. Farivar, Toward Application-Aware Security and Reliability. IEEE Security and Privacy, 2007. 5(1): p. 57-62.

[160] Pattabiraman, K., N. Nakka, Z. Kalbarczyk, and R. Iyer, Discovering Application-level Insider Attacks using Symbolic Execution, in Submission to 24th IFIP International Security Conference (SEC). 2008, IFIP: Cyprus, Greece.

[161] King, S.T., J. Tucek, A. Cozzie, C. Grier, W. Jiang, and Y. Zhou. Designing and implementing malicious hardware. 2008: USENIX Association Berkeley, CA, USA.

[162] Alkabani, Y.M. and F. Koushanfar, Active hardware metering for intellectual property protection and security.

[163] King, S.T., P.M. Chen, Y.-M. Wang, C. Verbowski, H. Wang, and J. Lorch. SubVirt: Implementing malware with virtual machines. in IEEE Symposium on Security and Privacy (Oakland). 2006.

[164] Rutkowska, J., Introducing Blue Pill. The official blog of the invisiblethings. org. June, 2006. 22.

[165] Garfinkel, T., K. Adams, A. Warfield, and J. Franklin. Compatibility is not transparency: VMM detection myths and realities.

[166] Upadhyaya, S., Real-Time Intrusion Detection with Emphasis on Insider Attacks. Lecture Notes in Computer Science, 2003. 2776: p. 82-85.

[167] Rhee, J., R. Riley, D. Xu, and X. Jiang, Defeating Dynamic Data Kernel Rootkit Attacks via VMM-based Guest-Transparent Monitoring.

[168] Baliga, A., V. Ganapathy, and L. Iftode. Automatic Inference and Enforcement of Kernel Data Structure Invariants. 2008: IEEE Computer Society Washington, DC, USA.

[169] Tanenbaum, A.S., J.N. Herder, and H. Bos, Can we make operating systems reliable and secure? Computer, 2006. 39(5): p. 44-51.

[170] Hieb, J. and J. Graham, Designing Security-Hardened Microkernels For Field Devices. Critical Infrastructure Protection II, 2008: p. 129.

[171] Akritidis, P., C. Cadar, C. Raiciu, M. Costa, and M. Castro, Preventing Memory Error Exploits with WIT, in Proceedings of the 2008 IEEE Symposium on Security and Privacy (sp 2008) - Volume 00. 2008, IEEE Computer Society.

[172] Barrantes, E.G., D.H. Ackley, S. Forrest, and D. Stefanovi, Randomized instruction set emulation. ACM Trans. Inf. Syst. Secur., 2005. 8(1): p. 3-40.

[173] Sovarel, A.N., D. Evans, and N. Paul, Where's the FEEB? the effectiveness of instruction set randomization, in Proceedings of the 14th conference on USENIX Security Symposium - Volume 14. 2005, USENIX Association: Baltimore, MD.

[174] Dean, D.W.a.R. Intrusion detection via static analysis. in International Conference on Security and Privacy (Oakland). 2001: IEEE.

[175] Giffin, J., Model-based intrusion detection system and evaluation, in Department of Computer Science. 2006, University of Wisconsin: Madison, WI.

[176] Chen, Y., R. Venkatesan, M. Cary, R. Pang, S. Sinha, and M.H. Jakubowski, Oblivious Hashing: A Stealthy Software Integrity Verification Primitive, in Revised Papers from the 5th International Workshop on Information Hiding. 2003, Springer-Verlag.

[177] Jacob, M., M.H. Jakubowski, and R. Venkatesan, Towards integral binary execution: implementing oblivious hashing using overlapped instruction encodings, in Proceedings of the 9th workshop on Multimedia \& security. 2007, ACM: Dallas, Texas, USA.

[178] Seshadri, A., M. Luk, E. Shi, A. Perrig, L.v. Doorn, and P. Khosla, Pioneer: verifying code integrity and enforcing untampered code execution on legacy systems, in Proceedings of the twentieth ACM symposium on Operating systems principles. 2005, ACM: Brighton, United Kingdom.

[179] Monrose, F., P. Wyckoff, and A.D. Rubin. Distributed Execution with Remote Audit. in ISOC Network and Distributed System Security Symposium. 1999.

[180] Provos, N., M. Friedl, and P. Honeyman, Preventing privilege escalation, in Proceedings of the 12th conference on USENIX Security Symposium - Volume 12. 2003, USENIX Association: Washington, DC.

[181] Brumley, D. and D. Song, Privtrans: Automatically partitioning programs for privilege separation.

[182] Pattabiraman, K., V. Grover, and B.G. Zorn, Samurai: protecting critical data in unsafe languages, in Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems 2008. 2008, ACM: Glasgow, Scotland UK.

[183] Nguyen-Tuong, A., D. Evans, J.C. Knight, B. Cox, and J.W. Davidson. Security through Redundant Data Diversity. in IEEE International Conference on Dependable Systems and Networks (DSN). 2008. Anchorage, AK: IEEE.

[184] Pattabiraman, K., Z. Kalbarczyk, and R. Iyer. Automated Derivation of Application-Aware Error Detectors using Static Analysis. in Internation Online Testing Symposium (IOLTS). 2007: IEEE.

[185] Abadi, M., M. Budiu, U.I. Erlingsson, and J. Ligatti. Control-flow integrity. in 12th ACM conference on Computer and communications security. 2005. Alexandria, VA, USA: ACM Press.

[186] Wu-ftp development group., WuFTP.

[187] NullLogic, NullHttpd web server. 2001.

[188] IBM. Linux TPM Device Driver. 2002 [cited; Available from: http://tpmdd.sourceforge.net/.

[189] Sen, K., D. Marinov, and G. Agha. CUTE: a concolic unit testing engine for C. in Proceedings of the 10th European software engineering conference held jointly with 13th ACM SIGSOFT international symposium on Foundations of software engineering. 2005. Lisbon, Portugal: ACM Press.

[190] Godefroid, P., N. Klarlund, and K. Sen. DART: directed automated random testing. in ACM SIGPLAN conference on Programming language design and implementation. 2005. Chicago, IL, USA: ACM Press.

[191] Bauer, L., J. Ligatti, and D. Walker. Composing security policies with polymer. in ACM SIGPLAN conference on Programming language design and implementation. 2005. Chicago, IL, USA: ACM Press.

[192] Rosu, G. and F. Chen. Towards Monitoring-Oriented Programming: A Paradigm Combining Specification and Implementation. in Third International Workshop on Runtime Verification. 2003. Boulder, Colorado: Elveiser.

[193] Gatzka, S. and C. Hochberger. On the Scope of Hardware Acceleration of Reconfigurable Processors in Mobile Devices. in 38th Annual Hawaii International Conference on System Sciences. 2005.

## **APPENDIX A: LIST OF PUBLICATIONS**

The following papers, technical reports and journal articles have been published based on

the work done in this dissertation. The chapters corresponding to the papers are indicated.

- Towards Application-aware Security and Reliability, Ravishankar Iyer, Zbigniew Kalbarczyk, Karthik Pattabiraman, William Healey, Peter Klemperer, Reza Farivar and Wen-Mei Hwu, *IEEE Security and Privacy Magazine*, January 2007. Material from this paper appears in Chapter 1.
- Application-based Metrics for Strategic Placement of Error Detectors, Karthik Pattabiraman, Zbigniew Kalbarczyk and Ravishankar Iyer, Proceedings of the International Symposium on *Pacific Rim Dependable Computing (PRDC)*, December 2005. Material from this paper appears in Chapter 2.
- Dynamic Derivation of Application-aware Error Detectors and their Hardware Implementation, Karthik Pattabiraman, Giacinto Paolo Sagesse, Daniel Chen, Zbigniew Kalbarczyk and Ravishankar Iyer, *Proceedings of European Dependable Computing Conference (EDCC)*, October 2006. Journal version submitted to the IEEE Transactions on Dependable and Secure Computing (TDSC). Material from this paper appears in Chapter 3.
- Automated Derivation of Application-aware Error Detectors using Static Analysis, Karthik Pattabiraman, Zbigniew Kalbarczyk and Ravishankar Iyer, Proceedings of *International Online Test Symposium (IOLTS)*, February 2007 -Journal version to appear in the IEEE Transactions on Dependable and Secure Computing (TDSC). Material from this paper appears in Chapter 4.
- SymPLFIED: Symbolic Program Level Fault-Injection and Error Detection Framework, Karthik Pattabiraman, Nithin Nakka, Zbigniew Kalbarczyk and Ravishankar Iyer, Proceedings of *International Conference on Dependable Systems and Networks (DSN), June 2008.* Material from this paper appears in Chapter 5.
- Discovering Application-level Insider Attacks using Symbolic Execution, Karthik Pattabiraman, Nithin Nakka, Zbigniew Kalbarczyk and Ravishankar Iyer *To appear at the IFIP International Information Security Conference (SEC), May 2009.* Material from this paper appears in Chapter 6.
- Insider Attack Detection by Information-Flow Signature Enforcement, with William Healey, Flore Yuan, Zbigniew Kalbarczyk and Ravishankar Iyer, Under submission. Material from this paper appears in Chapter 7.

## **AUTHOR'S BIOGRAPHY**

Karthik Pattabiraman received the Bachelor's degree in Information Technology from the University of Madras (India) in August 2001 and the Master's degree in Computer Science from the University of Illinois (Urbana-Champaign) in December 2004. His research focuses on the design of reliable and secure applications using innovations in compilers, formal methods and computer architecture. Based on his dissertation research, he was awarded the 2008 William C. Carter award by the IEEE Technical Committee on Fault-Tolerant Computing (TC-FTC) and the IFIP Working group on Dependability (WG 10.4) "to recognize an individual who has made a significant contribution to the field of dependable computing through his or her graduate dissertation research". Karthik's dissertation laid the foundation of the Trusted Illiac project at the University of Illinois, and Karthik was the lead student involved in building the Trusted Illiac prototype.

Karthik has done internships at Microsoft Research, IBM Research and Los Alamos National Labs, and has consulted for Microsoft Research. He has also been actively involved in the dependability community and co-organized the first and second workshops on Compiler and Architectural Techniques for Application Reliability and Security (CATARS) at the IEEE International Conference on Dependable Systems and Networks (DSN), 2008 and 2009. He is now a post-doctoral researcher at Microsoft Research and will soon join the University of British Columbia (UBC), Vancouver as an assistant professor of Electrical and Computer Engineering (ECE).