

Lecture Notes on Security Automata & Instrumentation

Matt Fredrikson

Carnegie Mellon University
Lecture 10

1 Introduction & Recap

We began last lecture by adding indirect control flow to our language, in the form of the instruction $\text{if}(Q) \text{ jump } e$. We then returned to our sandboxing approach for untrusted program α , and considered what could happen now that indirect jumps are in the language. Assuming that we are using SFI to enforce a sandboxing policy, there is still no way for α to read or write memory outside the sandbox. Is this true? Consider the following situation, where we can assume that SFI has been applied to the untrusted α starting at command 20.

```

      ⋮
10:  z := Mem(x)
11:  if(i ≥ 0) jump y
      ⋮
α { 20:  i := 0
    21:  x := attacker's desired address
    22:  y := 24
    23:  if(0 = 0) jump 10
    24:  copy memory contents from z
      ⋮
```

Here, our original program (not the untrusted α) dereferences memory and makes use of indirect control flow transfer. The attacker identifies a sequence of commands in the trusted portion of the program, and sets things up in a way so that unauthorized memory is copied into a variable that the attacker can later access once control is returned to the untrusted code. This should remind you of a *return-oriented programming* (ROP)

attacks [Sha07] that you learned about in 15-213. If we assume that the attacker knows the text of our program, then it is possible for them to identify “gadgets” in *code that we wrote* to do their bidding. But this crucially relies on the ability to change control flow using indirection so that commands are executed in the order needed by the attacker to carry out their goals.

We discussed addressing this with a very similar approach to what we did for SFI. We built a “code sandbox” between commands at pc_l and pc_h . Then we can rewrite indirect jump commands to ensure that their target always lies within these bounds.

$$\text{Rewrite all } \text{if}(Q) \text{ jump } e \text{ commands as } \text{if}(Q) \text{ jump } (e \& pc_h) | pc_l \quad (1)$$


This is a form of *control flow integrity* (CFI) [ABEL09], a technique for enforcing a broad class of safety properties that place limits on the allowed control flow paths in a program. But this sort of policy doesn’t give us much flexibility in terms of what kind of control flow we might want to restrict the untrusted code to, because any sequence within the code sandbox is considered valid.

We then looked at a more granular type of control flow policy that takes into account a control flow graph, and can be used to ensure that the target program executes only sequences of instructions corresponding to paths in the graph. Today we will generalize what we have learned about enforcing safety by introducing a type of policy that specifies safety using automata. This style of policy, called *security automata* [Sch00], is powerful enough to encapsulate all of the safety properties that can possibly be enforced at runtime, and is thus an indispensable tool for ensuring code safety.

2 Security automata


The point mentioned at the end of the previous section raises some interesting possibilities. What if we want to enforce a more general safety property that takes aspects of control flow and program state into account? For example, suppose that our language from the previous two sections has the ability to make three types of system calls, `send` and `recv` from the network and `read` from a local file. Then we quite naturally might want to enforce a safety policy on untrusted code which says that `send` cannot be called after `read`.

We can encode such a policy using a *security automaton* [Sch00]. Depicted below is a security automaton for the safety policy “no send after read”. The states are abstract in the sense that they do not reflect anything about the state of the program or what it is currently doing. Rather, they represent the state at which the policy is currently in. The transitions reflect facts about program state that must be true in order for the automaton to transition. In this case, `pc` denotes the current program counter, and `*pc` its contents. So for example `*pc ≠ read` corresponds to states in which the current instruction pointed to by the program counter is not read.



Notice that the only arrow going out of the rightmost state is a self-loop labeled $*pc \neq \text{send}$. There are no accepting states in a security automaton, and the way to interpret them is that as long as the automaton can transition from some arrow in its current state, then the policy has not been violated. So in this case, if the current policy state were the rightmost one, and the program entered into a state where $*pc = \text{send}$, then there would be no arrow to transition from and the policy would become violated.

Another way to think about it is that there is a “hidden” error state which corresponds to the policy being violated. Every node has a transition to the error state on the condition that is the negation of all other outgoing transitions from that state, as shown in the diagram below.



These definitions are equivalent, and we will continue using the convention that does not explicitly list the error state as this will reduce clutter in our diagrams.

Definition 1 (Security automaton[Sch00]). A security automaton is a nondeterministic state machine that consists of the following components:

- a countable set O of automaton states,
- a countable set $O_0 \subseteq O$ of initial states,
- a countable set Σ of transition symbols,
- a transition relation $\delta \subseteq O \times \wp(\Sigma) \times O$ between automaton states and sets of transition symbols.

We will assume that sets of automaton states are represented by formulas P, Q, \dots that can be evaluated on transition symbols to determine whether the symbol is in the set. Given a sequence of transition symbols $\sigma = \sigma_0, \sigma_2, \dots$, we say that the automaton accepts σ if and only if there is a corresponding sequence of states $o = o_0, o_1, \dots$ such that for each pair σ_i, σ_{i+1} in σ ,

- there is a corresponding pair o_i, o_{i+1} of states in o ,
- and there exists $(o_i, P, o_{i+1}) \in \delta$ where $P(\sigma_i)$ is true.

In other words, a trace is only accepted if there is a corresponding run of the automaton that always follows the transition function.

Definition 1 formally defines security automata in terms of a set of states O and transition symbols \mathcal{S} . We will generally assume that \mathcal{S} is the set of all program states, so that we can describe program traces as being accepted or not by a security automaton. This also implies that sets of states in the transition relation are defined in terms of formulas on program states, which we have already studied extensively.

2.1 Enforcing security automata policies

The primary means of enforcing policies defined using security automata is with a *reference monitor* (RM). The RM is a mechanism that examines the program as it executes, using information about the current and past states to decide whether the policy has been violated. This is done according to Definition 2, and was sketched out at the beginning of this section.

Definition 2 (Security automaton enforcement). Let O_c be the current set of states that the security automaton is in. Then for each step that the program is about to take resulting in new program state ω , the reference monitor does one of two things.

1. For each state $o \in O_c$ that the automaton can transition from, the states $\delta(o, P, o')$ for all transition edges where $P(\omega)$ is true are added to the new automaton states.
2. If the automaton cannot transition from any of its current states, then the program is not allowed to enter state ω and the reference monitor takes remedial action.

As long as the policy is not violated, then the RM allows the program to continue executing as it otherwise would. If the policy is violated, then the RM intervenes on the program execution to take some remedial action. This could mean simply aborting the execution, or something less drastic that prevents harm in other ways.

Necessary assumptions. As pointed out by Schneider in his seminal work on security automata [Sch00], there are several assumptions that one must make in order to enforce these policies effectively with a reference monitor. First, the reference monitor needs to simulate the execution of the automaton as the program runs, so it must keep track of which state the policy is in on the actual hardware running the program. This means that the automaton cannot require an unbounded amount of memory, so automata that have an infinite number of states are not in general enforceable.

Second, the RM must be able to prevent the program from entering a state that would result in a policy violation. This is called *target control*, and is a more subtle issue that it may at first seem. Take for example the policy of “real-time” availability, which states

that a principal should not be denied a resource for more than n real-time seconds. How could a reference monitor enforce this policy? It might try to predict the amount of time that it takes to remediate a trace that is about to violate the policy, and take action earlier than necessary to prevent the violation. But how does it know that the policy would have actually been violated in this case? Unless the reference monitor can literally stop time, this is not an enforceable policy.

Third, the program under enforcement must not be able to intervene directly on the state of the reference monitor. This is called *enforcement mechanism integrity*, and is crucial for ensuring that the policy defined by the automaton is the one that is actually enforced on the target program. We dealt with an instance of this issue earlier in the lecture, when we used control flow integrity to make sure that inlined safety checks weren't bypassed by indirect jumps. But now that the policy itself has state, the enforcement mechanism must also guarantee that the target program does not make changes to that state or influence it in any way that doesn't follow the automaton transitions.

Inline SA enforcement. One approach to implementing security automata enforcement uses inlined checks to update and maintain state set aside to simulate the automaton. If we assume that formulas on SA transitions are formulas over program states, and there are N security automata states, then we can set aside a region of N memory cells at addresses a_{sa} through $a_{sa} + N$ to hold the current state of the automaton. If $\text{Mem}(a_{sa} + i)$ is non-zero, then we assume that the automaton has entered into state i , and otherwise not.

Next we need to implement the transition function, updating the contents of $\text{Mem}(a_{sa}) - \text{Mem}(a_{sa} + N)$ to simulate the automaton. Suppose that the automaton has an edge from states i to j labeled with formula P . Then for each instruction in the program we compute the verification condition of (2).

$$[\alpha]\neg P \tag{2}$$

If (2) is valid before executing α , then it means that all traces after executing α will satisfy $\neg P$. On the other hand, if it is not valid, then at least one trace of α may satisfy P . This means that we need to insert a check whenever Eq 2 is not valid.

What check do we insert? At runtime, we will be in a particular state ω . We want to know if after executing α , P will be true, and if it is, then update the state of the automaton. We can accomplish this by simply checking that $\omega \models [\alpha]P$. Of course, we will want to use axioms to remove the box modality so that the check is actually in terms of arithmetic, and can be easily evaluated.

So we insert instrumentation immediately before α that checks $\text{Mem}(a_{sa} + i) \neq 0 \wedge [\alpha]P$, and if it is true then sets $\text{Mem}(a_{sa} + j)$ to a non-zero value. Then for each state i in the SA, we compute similar checks for transition to the "error state". If i has outgoing edges labeled P_1, \dots, P_n , we insert a check for:

$$\text{Mem}(a_{sa} + i) \neq 0 \wedge [\alpha]\neg(P_1 \vee \dots \vee P_n) \tag{3}$$

If this check passes, it means that the automaton cannot transition from state i . If this holds for every state in the automaton, then the instrumentation aborts execution.

The instrumentation described so far only addresses updates to the SA state. We must also take steps to ensure the integrity of the inlined mechanism, and there are two sources of vulnerability.

- The contents of $\text{Mem}(a_{sa}) - \text{Mem}(a_{sa} + N)$ must not be modified by any part of the program except the inserted instrumentation. Applying software fault isolation to the untrusted instructions can ensure that this aspect of integrity holds.
- The inserted instrumentation could be subverted by indirect control flow. Enforcing CFI on the untrusted code using the original control flow graph ensures that this will not happen.

This is sufficient to implement a basic inlined security automaton enforcement mechanism. However, it may impose a severe performance overhead due to all the safety checks.

3 Dynamic instrumentation

We have been discussing policy enforcement in a somewhat idealized model, where we assume that programs are given to us as source code in a simple language with few instructions. In the “real world” this is not usually the case, and we may be forced to deal with large untrusted programs given to us to execute at runtime, and possibly without source code. So we must find a way to enforce policies on bytecode, and presumably fast lest we introduce unacceptable latency into the system.

Suppose that we wish to implement the inline security automata enforcement scheme from the previous section by changing the instructions throughout the program prior to running it. This seems like a reasonable approach, because the scheme just requires that we check verification conditions on each instruction and replace them when necessary. All that we need to assume is the ability to identify instructions, and compute verification conditions.

3.1 Challenges for static instrumentation

But bytecode programs on modern architectures like x86 and AMD64/Intel 64 are extremely difficult to reason about statically, and it may not even be possible to identify which instructions the program will end up executing. One practical issue is the fact that programs can generate new instructions by writing to memory, and then use an indirect jump to begin executing the newly-written code. This can be mitigated by the operating system with a *Write XOR Execute* policy, which ensures that any page of memory may be either writeable or executable, but not both. This is effective, but makes some functionality extremely difficult to implement such as language interpreters that do on-the-fly compilation and optimization.

Even with *Write XOR Execute*, the presence of indirect control flow and variable-length instruction encoding makes it impossible to tell which instructions will actually be executed. The program can do an arbitrarily complicated computation to derive a

target address in existing code, so that the static analysis is unable to determine where execution will resume after a jump. If the target address is in the middle of an existing instruction, it may result in a completely different program being executed. Consider the following example, taken from [Sha07].

<i>Bytecode</i>	<i>Instruction</i>	
f7 c7 07 00 00 00	test \$0x00000007, %edi	(4)
0f 95 45 c3	setnzb -61(%ebp)	

This code is taken from the entry point of an encryption routine in the GNU C library, often referred to as simply `libc`. If execution begins one byte after the entry point of (4), a completely different program is executed.

<i>Bytecode</i>	<i>Instruction</i>	
c7 07 00 00 00 0f	movl \$0x0f000000, (%edi)	(5)
95	xchg %ebp, %eax	
45	inc %ebp	
c3	ret	

Importantly this implies that given a sequence of bytecodes, there are numerous possible programs that could end up being executed depending on which addresses are targeted by indirect jumps. In order to instrument the right one, a static analysis needs to determine what these addresses will be, and this is an undecidable problem in general. Moreover, it could be that information not available statically, such as network packets, are used in part to compute target addresses, adding yet another very plausible complication for static instrumentation in this setting.

3.2 Instrumenting with just-in-time compilation

Perhaps a better approach given these challenges is to delay “code discovery” until the program is actually running. This is helpful for many reasons.

- If the program generated instructions in memory and transferred control to them, we no longer need to infer what those instructions will be. We can simply wait until the program has already written them, and instrument them immediately before the control transfer.
- If a program executes an indirect jump, we do not need to predict what the target address will be. We simply wait until immediately before the jump is executed, at which point the target address will be stored in memory or a register, and begin instrumenting the target of the jump.
- Some other cases that we have not discussed are handled similarly, such as libraries that are loaded after the program begins executing. In each such case, the instrumentation is delayed until immediately before the instructions in question begin executing, at which point all of the necessary information is available.

The obvious drawback to this approach is the fact that we need to examine the execution as it unfolds, rewriting instructions whenever necessary as dictated by the policy.

Just-in-time compilation. A successful and widely-deployed approach to mitigate the performance penalty imposed by such a scheme is called *just-in-time (JIT) compilation* [Ayc03]. The key insight behind JIT compilation is to increase the granularity at which the instrumenter examines code at runtime, looking at “chunks” of instructions rather than individual ones.

Increasing the granularity in this way allows the instrumenter to compile instruction chunks, with their instrumentation included, on the fly into optimized code that is then executed directly. Further performance enhancements can then be layered on top of this basic approach, such as caching previously-compiled chunks to save redundant work, as well as more aggressive optimizations to sequences of chunks that end up being executed more often.

The question then becomes what constitutes a chunk. Larger chunks will generally create more opportunities for optimization, and because more of the instructions are dealt with each time, require fewer (expensive) calls to the compiler. However, this tendency is limited by the fact that if a chunk crosses an indirect control flow instruction, then we run into exactly the same problems we are trying to avoid with dynamic instrumentation in the first place. Even if our chunks cross direct, predictable control flow branches, then we run the risk of doing unnecessary compilation and instrumentation by processing multiple branches when the execution will only end up following one of them.

The typical approach is to use *basic blocks* as chunks. A basic block is a contiguous sequence of instructions that ends in a control flow transfer instruction (e.g., `jmp`, `ret`, `call`, ...). For example, the sequence of instructions in (5) is a basic block because it ends with a `ret` instruction, which transfers control to the instruction pointed to by the return address on the stack. On the other hand, (4) is not a basic block because it does not end in such an instruction.

Using basic blocks as chunks, the instrumenter will begin scanning a sequence of bytecodes until it reaches a control transfer instruction. It will then instrument each of the instructions in the basic block as prescribed by the policy, compile the resulting instructions, and execute them. However, it must ensure that it regains control when the basic block is finished executing. It then begins scanning instructions again at the bytecodes pointed to by the instruction pointer, repeating the process all over again. In this way we can be sure that exactly the code that is executed is instrumented according to the policy.

A look ahead: Pin. In your next lab, you will make use of an instrumentation tool called Pin [LCM+05] that is based on JIT compilation. Pin is under ongoing development by Intel, and is widely used in industry as well as in academic research. It simplifies the task of instrumenting binaries at runtime by providing a high-level API for both inspecting and instrumenting sequences of instructions at runtime.

To see why this is helpful, consider the task of instrumenting an x86 binary to prevent writes to certain portions of memory. To do so, we must rewrite all instructions that can change memory with instrumentation to stop the unwanted writes. Which instructions

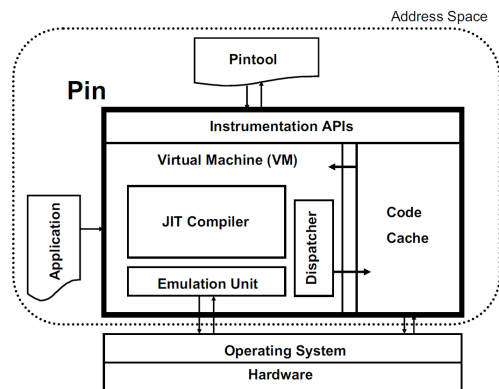


Figure 1: Pin software architecture (from [LCM⁺05]).

can change memory? The obvious ones are `mov`, `push`, `pop`, `lea`, `xchg`, and perhaps a few others. But what about the many variants of `mov`, such as `movsb`, `movsw`, `movz`, `movzx`? Do the other instructions have variants as well, and how can we be sure that we've covered each one? Pin simplifies things for us by providing `INS_IsMemoryWrite(ins)`, which returns true if `ins` can update memory.

Figure 1 shows the architecture of Pin. Users interact with it by writing a "Pintool", which is a conventional C or C++ program that makes use of the Pin inspection and instrumentation API. To run a compiled program under the pintool's instrumentation, the program's binary is passed to Pin along with the compiled pintool. Pin then takes care of just-in-time compiling the target program, and can invoke callbacks to the pintool as requested for inspection, or rewrite instructions as requested for instrumentation. As execution proceeds, Pin's optimization routines run in tandem to progressively optimize the compiled code.

You will learn more about the specifics of the Pin API in the handout for the next lab, and get hands-on experience using it to implement SFI as well as a security automaton policy. For more detailed information on how Pin works, consult the original paper [LCM⁺05].

References

- [ABEL09] Martín Abadi, Mihai Budiu, Úlfar Erlingsson, and Jay Ligatti. Control-flow integrity: Principles, implementations, and applications. *ACM Transactions on Information and Systems Security*, 13(1):4:1–4:40, November 2009.
- [Ayc03] John Aycock. A brief history of just-in-time. *ACM Computing Surveys*, 35(2):97–113, June 2003.
- [LCM⁺05] Chi-Keung Luk, Robert Cohn, Robert Muth, Harish Patil, Artur Klauser, Geoff Lowney, Steven Wallace, Vijay Janapa Reddi, and Kim Hazelwood.

- Pin: Building customized program analysis tools with dynamic instrumentation. In *Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 2005.
- [Sch00] Fred B. Schneider. Enforceable security policies. *ACM Transactions on Information Systems Secur.*, 3(1):30–50, February 2000.
- [Sha07] Hovav Shacham. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In *Proceedings of CCS 2007*, October 2007.