

Computer Sciences Department

Secure Programming via Visibly Pushdown Safety Games

William Harris

Somesh Jha

Thomas Reps

Technical Report #1710

January 2012



WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON

Secure Programming via Visibly Pushdown Safety Games

William R. Harris¹, Somesh Jha¹, and Thomas Reps^{1,2}

¹ University of Wisconsin-Madison, Madison, WI, USA
{wrharris, jha, reps}@cs.wisc.edu

² GrammaTech, Inc., Ithaca NY, USA

Abstract. Several recent operating systems provide system calls that allow an application to explicitly manage the privileges of modules with which the application interacts. Such *privilege-aware operating systems* allow a programmer to write a program that satisfies a strong security policy, even when it interacts with untrusted modules. However, it is often non-trivial to rewrite a program to correctly use the system calls to satisfy a high-level security policy. This paper concerns the *policy-weaving problem*, which is to take as input a program, a desired high-level policy for the program, and a description of how system calls affect privilege, and automatically rewrite the program to invoke the system calls so that it satisfies the policy. We present an algorithm that solves the policy-weaving problem by reducing it to finding a winning modular strategy to a visibly pushdown safety game, and applies a novel game-solving algorithm to the resulting game. Our experiments demonstrate that our algorithm can efficiently rewrite practical programs for a practical privilege-aware system.

1 Introduction

Developing practical but secure programs remains a difficult, important, and open problem. Web servers and VPN clients execute unsafe code, and yet are directly exposed to potentially malicious inputs from a network connection [?]. System utilities such as Norton Antivirus scanner [?], `tcpdump`, the DHCP client `dhclient` [?], and file utilities such as `bzip`, `gzip`, and `tar` [?,?,?] have contained unsafe code with well-known vulnerabilities that allow them to be compromised if an attacker can control their inputs. Once an attacker compromises any of the above programs, they can typically perform any action allowed for the user that invoked the program, because the program does not restrict the privileges with which its code executes.

Traditional operating systems provide to programs only weak primitives for managing their privileges [?,?,?]. As a result, if a programmer is to verify that his program is secure, he typically must first verify that the program satisfies very strong properties, such as memory safety. However, recent work [?,?,?] has produced new operating systems that allow programmers to develop programs that execute unsafe code but still satisfy strong properties, and to construct such programs with significantly less effort than fully verifying the program. Such systems map each program to a set of privileges, and extend the set of system calls provided by a traditional operating system with security-specific

calls (which henceforth we will call *security primitives*) that the program invokes to manage its privileges. We call such systems *privilege-aware systems*.

This paper concerns the *policy-weaving problem*, which is to take a program and a security policy that defines what privileges the program must have, and to automatically rewrite the program to correctly invoke the primitives of a privilege-aware system so that the program satisfies the policy when run on the system. The paper addresses two key challenges that arise in solving the policy-weaving problem. First, a privilege-aware system cannot allow a program to modify its privileges arbitrarily, or an untrusted module of the program could simply give itself the privileges that it requires to carry out an attack. Instead, the system allows a program to modify its privileges subject to system-specific rules. In practice, these rules are subtle and difficult to master; the developers of the Capsicum capability system reported issues in rewriting the `tcpdump` network utility to use the Capsicum primitives to satisfy a security policy, while preserving the original functionality of `tcpdump` [?].

Second, the notions of privilege often differ between privilege-aware systems, and thus so too do the primitives provided by each system, along with the rules relating privileges to primitives. The Capsicum operating system defines privileges as capabilities [?], the Decentralized Information Flow Control (DIFC) operating systems Asbestos, HiStar, and Flume [?,?,?] define privileges as the right to send information, and each provide different primitives for manipulating information-flow labels [?]. Thus, a policy-weaving algorithm for a specific system must depend on the privileges and primitives of the system, yet it is undesirable to manually construct a new policy-weaving algorithm for each privilege-aware system that has been or will yet be developed.

We address the above challenges by reducing the policy-weaving problem to finding a winning Defender strategy to a two-player safety game. Each game is played by an Attacker, who plays program instructions, and a Defender, who plays system primitives. The game accepts all sequences of instructions and primitives that violate the given policy. A winning Defender strategy never allows the Attacker to generate a play accepted by the game, and thus corresponds to a correct instrumentation of the program, which invokes primitives so that the policy is never violated. If the rules describing how a system's primitives modify privileges can be encoded as an appropriate automaton, then the game-solving algorithm can be applied to rewrite programs for the system. We argue that stack-based automata, in particular *visibly pushdown automata* (VPAs) [?], are sufficient to model the rules of practical privilege-aware systems. Furthermore, *modular* winning strategies exactly correspond to correct instrumentations of programs for such systems.

Finding a modular winning strategy to a game defined by a VPA is NP-complete. However, games resulting from policy-weaving problems are constructed as products of input automata, and a game will often have a strategy whose structure closely matches one of the inputs. Inspired by this observation, we present a novel algorithm that, given a game, finds a modular strategy with structure similar to an additional, potentially smaller game called a *scaffold*. We show that our scaffolding algorithm generalizes two known algorithms for finding modular

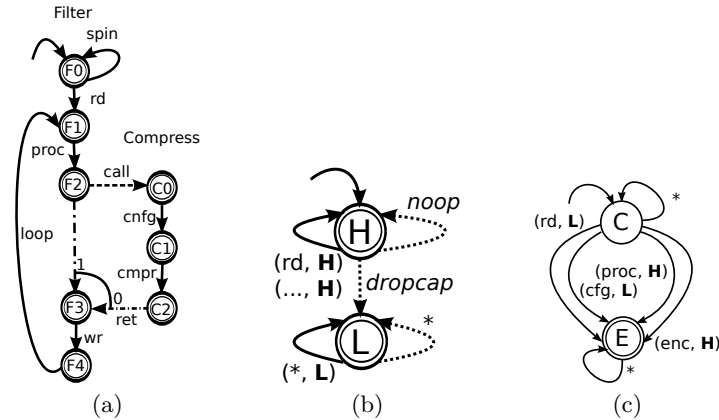


Fig. 1. Automata models of (a) the program `Filter`, (b) `Filter`'s MiniCap monitor, and (c) `Filter`'s policy for MiniCap. Notation is explained in §??.

strategies [?,?] — in particular, those algorithms result from using two (different) “degenerate” scaffolds, and correspond to two ends of a spectrum of algorithms that can be implemented by our algorithm. We evaluated the scaffold-based algorithm on games corresponding to policy-weaving problems for six UNIX utilities with known vulnerabilities for the Capsicum capability system, and found that it could rewrite practical programs efficiently, and that the choice in scaffold often significantly affected the performance of the algorithm.

Organization §?? motivates by example the policy-weaving problem and our game-solving algorithm. §?? defines the policy-weaving problem, and reduces the problem to solving visibly-pushdown safety games. §?? presents a novel algorithm for solving visibly pushdown safety games. §?? presents an experimental evaluation of our algorithm. §?? discusses related work.

2 Overview

In this section, we motivate the policy-weaving problem. We sketch how the policy-weaving problem can be reduced to finding a winning strategy to a class of safety games, and how the structure of games constructed from policy-weaving problems makes them amenable to our novel game-solving algorithm.

2.1 An Example Policy-Weaving Problem: `Filter` on MiniCap

We illustrate the policy-weaving problem using an example program `Filter` that reads information from an input channel, processes and compresses the data, and then writes the data to an output channel. `Filter` is inspired by the UNIX utilities `tcpdump` and `gzip`, which have exhibited security vulnerabilities in the past, and have previously been rewritten manually for the Capsicum privilege-aware systems [?]. The executions of `Filter` are presented as the runs of the automaton F in Fig. ??(a), where each transition is labeled with a program action. Intraprocedural transitions are denoted with solid lines. Call transitions, which place their source node on a stack (see Defn. ??), are denoted with dashed lines. Return transitions, defined by the top two states of the stack, are denoted

with dash-dot lines, where the transition from the top state of the stack is labeled with a 0, and the transition from the next state down on the stack is labeled with a 1. In each figure, doubled circles denote accepting states.

Filter executes a no-op instruction (**spin**) until it reads data from its designated input channel (e.g., UNIX **stdin**) (**read**), processes a segment of its input data (**proc**), and calls a compression function **Compress** (**call**). **Compress** first opens and reads a configuration file (**cnfg**), compresses its input data (**cmpr**), and returns the result (**ret**). After **Compress** returns, **Filter** writes the data to its designated output channel (e.g., UNIX **stdout**) (**wr**), and loops to read another segment of data (**loop**).

Unfortunately, in practice, much of the code executed by a practical implementation of functions like **Filter** and **Compress** (e.g., **tcpdump** and **gzip** [?]) is not memory-safe, and thus allows an attacker to violate the security policy of a program. Suppose that the programmer wants to ensure that **Filter** only interacts with communication channels by opening and reading from its designated input at **read** and writing to its designated output at **wr**, and **Compress** only interacts with communication channels by reading from its configuration files at **cnfg**. However, suppose also that the data-processing action **proc** in **Filter** and the compression action **cmpr** in **Compress** perform memory-unsafe operations when passed particular inputs. Then an attacker who can control the inputs to **Filter** could craft a malicious input that injects code that opens a communication channel (e.g., a file) and violates the policy.

However, if the programmer correctly rewrites **Filter** for a suitable *privilege-aware systems*, then the rewritten **Filter** will satisfy such a policy even if it executes code injected by an attacker. Consider a privilege-aware system MiniCap, which is a simplification of the Capsicum capability system now included in the “RELEASE” branch of FreeBSD [?,?]. MiniCap maps each executing process to a two-valued flag denoting if the process has high or low privilege. If a process has high privilege H, then it can open communication channels, but if it has low privilege L, then it can only read and write to its designated input and output channels. A process on MiniCap begins executing with high privilege, but may invoke the MiniCap primitive **dropcap**, which directs MiniCap to give the process low privilege, and never give the process the high privilege again. A process thus might invoke **dropcap** after executing safe code that requires high privilege, but before executing unsafe code that requires only low privilege.

MiniCap also allows one process to communicate with another process via a *remote procedure call (RPC)*, in which case the called process begins execution with high privilege, independent of the privilege of the caller. The Capsicum capability system uses RPC in this way, while DIFC systems allow a process to call a process with different privileges via an analogous *gate call* [?].

MiniCap is partially depicted in Fig. ??(b) as an automaton **M** that accepts sequences of privilege-instruction pairs and primitives executed by **Filter**. We call **M** the MiniCap *monitor* of **Filter**. **M** accepts a trace of privilege-instruction pairs and primitives if when **Filter** executes the sequence of instructions and primitives, MiniCap grants **Filter** the privilege paired with each instruction. The call and return transitions of **M** are omitted for simplicity; **M** transitions on

an RPC to the high-privilege state H, and returns from an RPC to the calling state.

Filter's policy can be expressed directly in terms of MiniCap's privileges by requiring that the instructions `read` and `cnfg` execute with high privilege, while the instructions `proc` and `cmpr` execute with low privilege. The policy is presented as an automaton `Pol` in Fig. ??(c), where each transition is labeled with a privilege-instruction pair (the label * denotes any label that does not appear explicitly on a transition from the same source state). The traces accepted by `Pol` are the sequences of instruction-privilege pairs that violate the policy.

For **Filter** to satisfy its policy when it is run on MiniCap, it must use the primitives of MiniCap in a way that is only indirectly related to, and significantly more complex than, its desired policy. (1) invoke `dropcap` after executing `read` but before executing `proc`, (2) call `Compress` via RPC so that `Compress` executes `cnfg` with high privilege, (3) invoke `dropcap` after executing `cnfg` but before executing `cmpr`. This rewritten **Filter** is "modular" across calls and returns, in the sense that the rewritten **Filter** and `Compress` invoke primitives independently of the actions of each other. On practical privilege-aware systems, a process that can be called via RPC cannot necessarily trust its caller, and thus cannot trust information passed by its caller. Thus a practical instrumentation must be modular.

The policy-weaving problem for **Filter** is to take `F`, its policy `Pol`, and MiniCap monitor `M`, and instrument **Filter** to use MiniCap's primitives modularly to satisfy `Pol`.

2.2 Policy-Weaving Filter via Safety Games

Each policy-weaving problem can be reduced to finding a winning strategy to a safety game. A safety game is played by two players, an Attacker and Defender, and is a transition system in which each state belongs to either the Attacker or the Defender. The goal of the Attacker is to drive the state of the game to an accepting state, while the goal of the Defender is to thwart the Attacker. The game is played in turns: when the game enters an Attacker state, the Attacker chooses the next transition, and when the game enters a Defender state, the Defender chooses the next transition. A strategy for the Defender takes as input a play of the game, and chooses the next transition for the Defender. A winning strategy chooses Defender transitions that never allow the Attacker to drive the state of the game to an accepting state.

From program, policy, and monitor automata, we can construct a game that accepts all policy-violating executions of a version of the program that is instrumented to invoke the primitives of the monitor. The game is constructed by (1) transforming the alphabets of the automata to a common alphabet defined by the instructions, privileges, and primitives, (2) constructing the product of the transformed automata, and (3) transforming the alphabet of the resulting product game so that all Attacker transitions are labeled with program instructions, and all Defender transitions are labeled with system primitives.

A subset of the transitions of the game G_{ex} constructed from `F`, `Pol`, and `M` are shown in Fig. ??(a). Each state of G_{ex} is either an Attacker or Defender state

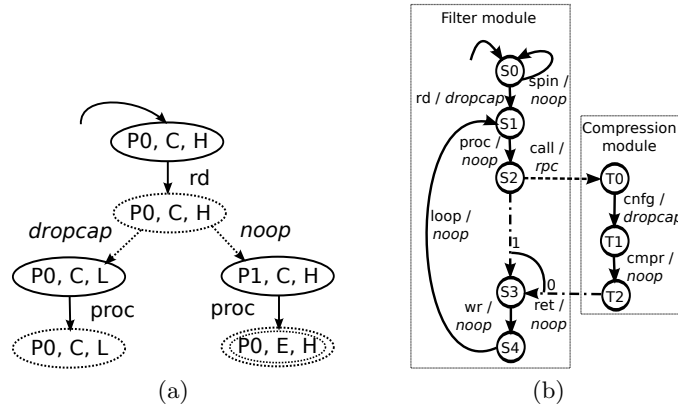


Fig. 2. (a) a selection of transitions of the game G_{ex} that is the product of F , Pol , and M ; (b) a strategy corresponding to a correct instrumentation of **Filter**.

constructed from a triple of a state of F , state of Pol , and state of M , and each state in Fig. ??(a) is labeled with its triple. Each Attacker state and Attacker transition is denoted with a solid circle or line, while each Defender state is denoted with a dotted circle or line. The play “read, noop, proc” is accepted by G_{ex} (i.e., is a winning play for the Attacker) because it is an execution in which the instrumented **Filter** does not execute **dropcap** before executing **proc**, causing **proc** to execute with high-privilege, which violates the policy Pol . However, the play “read, dropcap, proc” is not accepted by G_{ex} , because it corresponds to an execution in which **Filter** invokes **dropcap**, causing **proc** to execute with low privilege, which satisfies the policy.

One winning Defender strategy to G_{ex} , which corresponds to the correct instrumentation of **Filter** given in §??, is presented in Fig. ??(b). The strategy is a transducer that, from its current state, reads an instruction executed by **Filter**, outputs the primitive paired with the instruction on the label of a transition t (Fig. ??(b) includes a primitive **noop** that denotes that no MiniCap primitive is invoked), transitions on t , and reads the next instruction. The strategy is partitioned into a **Filter** module that chooses what primitives are invoked during an execution of **Filter**, and a **Compress** module that chooses primitives are invoked during the execution of **Compress**. The modules are independent, in that the primitives chosen by the **Compress** module are independent of the instructions and primitives executed by **Filter** before the most recent call of **Compress**.

Solving games constructed from policy-weaving problems efficiently is a hard problem. The game G_{ex} is the product of F , Pol , and M , and thus has a state space whose size is proportional to the product of the sizes of F , Pol , and M (G_{ex} has 128 states). Furthermore, finding modular winning Defender strategies to games is NP-complete in the size of the game. However, in practice, games constructed from policy-weaving problems have a winning strategy whose structure closely matches the structure of one of the input automata. For example, the winning strategy in Fig. ??(b) closely matches the structure of F . Each execution of

Filter is in state F_n of F when the strategy is in state S_n , and in state C_n of F when the strategy is in state T_n (see Fig. ??(a) and Fig. ??(b)). To find winning modular strategies to games efficiently, we apply a novel algorithm that takes a game and an additional, potentially smaller, game called a *scaffold*, and searches for a winning strategy whose structure is similar to that of the scaffold. For G_{ex} , F serves as such a scaffold.

3 Policy Weaving as a Safety Game

3.1 Definition of the Policy-Weaving Problem

The policy-weaving problem is to take a program, a description of a privilege-aware system, and a policy that describes what privileges the program must have as it executes on the system, and to instrument the program so that it always has the privileges required by the policy. We model a program, policy, and privilege-aware system each as a Visibly Pushdown Automaton.

Definition 1. A *deterministic visibly-pushdown automaton* (VPA) for internal actions Σ_I , call actions Σ_C , and return actions Σ_R (alternatively, a $(\Sigma_I, \Sigma_C, \Sigma_R)$ -VPA) is a tuple $V = (Q, \iota, Q_F, \tau_i, \tau_c, \tau_r)$, where: Q is the set of *states*; $\iota \in Q$ is the *initial state*; $Q_F \subseteq Q$ is the set of *accepting states*; $\tau_i : Q \times \Sigma_i \rightarrow Q$ is the *internal transition function*; $\tau_c : Q \times \Sigma_c \rightarrow Q$ is the *call transition function*; $\tau_r : Q \times \Sigma_r \times Q \rightarrow Q$ is the *return transition function*.

For $\widehat{\Sigma} = \Sigma_I \cup \Sigma_C \cup \Sigma_R$, each VPA accepts a set of *traces* of (i.e., sequences of actions in) $\widehat{\Sigma}$. Let ϵ denote the empty sequence. Let “.” denote the concatenation of two sequences; for set X , $x \in X$, and $s \in X^*$, $x \cdot s = [x] \cdot s$ and $s \cdot x = s \cdot [x]$, where $[x] \in X^*$ is the sequence containing only x . For sets X_0 and X_1 , let $X_0 \cdot X_1$ be the set of all sequences $x_0 \cdot x_1$ for $x_0 \in X_0$ and $x_1 \in X_1$. Let $\tau : Q^* \times \widehat{\Sigma} \rightarrow Q^*$ map a sequence of states $s \in Q$ (i.e., a *stack*) and action $a \in \widehat{\Sigma}$ to the stack to which V transitions from s on a :

$$\begin{aligned} \tau(q \cdot s, a) &= \tau_I(q, a) \cdot s && \text{for } a \in \Sigma_I \\ \tau(q \cdot s, a) &= (\tau_C(q, a) \cdot q \cdot s) && \text{for } a \in \Sigma_C \\ \tau((q \cdot s_0 \cdot s'), a) &= (\tau_R(q, a, s_0), s') && \text{for } a \in \Sigma_R \end{aligned}$$

Let $\rho : \widehat{\Sigma}^* \rightarrow Q^*$ map each trace to the stack that V is in after reading the trace. Formally, $\rho(\epsilon) = \iota$, and for $a \in \widehat{\Sigma}$ and $s \in \widehat{\Sigma}^*$, $\rho(s \cdot a) = \tau(\rho(s), a)$. A trace $t \in \widehat{\Sigma}^*$ is accepted by V if $\rho(t) = (q, s)$ with $q \in Q_F$. In a trace t , an instance c of a call action is *matched* by an instance r of a return action if c is before r in t , and each instance c' of a call action in t between c and r is matched by a return action r' between c and r . A trace is *matched* if all call and return actions in the string are matched. Let $\mathcal{L}(V)$ be the set of all traces accepted by V . \square

A program is a language of traces of intraprocedural instructions, calls, and returns of the program (e.g., for **Filter** in §??, **spin**, **read**, etc.). Let $\text{Instrs} = (\Sigma_I, \Sigma_C, \Sigma_R)$, let $\widehat{\text{Instrs}} = \Sigma_I \cup \Sigma_C \cup \Sigma_R$, and let a program P be an Instrs -VPA.

A program policy is a language of traces of program instructions paired with privileges. A program’s privilege is a system-specific ability (e.g., for MiniCap

in §??, a program may have either the high privilege H or the low privilege L). Let Privs be a set of privileges, and let the set of *privileged executions* of P be $(\widehat{\text{Instrs}} \times \text{Privs})^*$. Let an $(\text{Instrs}, \text{Privs})$ -*policy* for P be a $(\Sigma_I \times \text{Privs}, \Sigma_C \times \text{Privs}, \Sigma_R, \times \text{Privs})$ -VPA (e.g., Fig. ??(c)) that accepts all privileged executions that constitute violations.

A privilege-aware monitor is a language of privileged executions interleaved with primitives. The primitives of a privilege-aware system are the set of security-specific system calls that the application can invoke to manage its privileges (e.g., for MiniCap, the system call `dropcap`). Let Prims be a set of primitives and let the *instrumented executions* of P be $(\text{Prims} \cdot \text{Instrs})^*$. A privilege-aware monitor of P reads an instrumented execution of P, and decides what privilege P has as it executes each instruction. Let an $(\text{Instrs}, \text{Privs}, \text{Prims})$ -privilege-aware monitor M be a $((\Sigma_I \times \text{Privs}) \cup \text{Prims}, \Sigma_C \times \text{Privs}, \Sigma_R \times \text{Privs})$ -VPA.

Definition 2. (Policy-Weaving Problem) Let P be a program with internal, call, and return alphabets $\text{Instrs} = (\Sigma_I, \Sigma_C, \Sigma_R)$. For privileges Privs , let Pol be an $(\text{Instrs}, \text{Privs})$ -policy of P. For primitives Prims , let M be an $(\text{Instrs}, \text{Privs}, \text{Prims})$ -privilege-aware monitor.

Let an *instrumentation function* be a function $I : \text{Instrs}^* \rightarrow \text{Prims}$, and let $I_{\text{tr}} : \text{Instrs}^* \rightarrow (\text{Prims} \text{ Instrs})^*$ map each sequence of instructions to the instrumentation of the sequence defined by I : $I_{\text{tr}}(\epsilon) = I(\epsilon)$, and for $s \in \text{Instrs}^*$ and $a \in \text{Instrs}$, $I_{\text{tr}}(s \cdot a) = I_{\text{tr}}(s) \cdot I(s \cdot a) \cdot a$. Let $\text{PrivEx}_M : (\text{Prims} \text{ Instrs})^* \rightarrow (\text{Instrs} \times \text{Privs})^*$ map each instrumented execution to the privileged execution that it induces on M: for primitives p_j , instructions i_j , and privileges r_j , $\text{PrivEx}_M([p_0, i_0, \dots, p_n, i_n]) = [(p_0, r_0), \dots, (p_n, r_n)]$ if and only if $[p_0, (i_0, r_0), \dots, p_n, (i_n, r_n)] \in \mathcal{L}(M)$.

The *policy-weaving problem* $\text{POLWEAVE}(P, \text{Pol}, M)$ is to find an instrumentation function I such that:

1. I instruments P to never violate the Pol: $\text{PrivEx}_M(I_{\text{tr}}(\mathcal{L}(P))) \cap \mathcal{L}(\text{Pol}) = \emptyset$.
2. I chooses primitives independently of the execution before the most recent call; i.e., I is *modular*. Let $p^0, p^1 \in \text{Img}(I_{\text{tr}})$, (where for a relation R , $\text{Img}(R)$ is the image of R), and $p^0 = p_0^0 \cdot c \cdot p_1^0 \cdot r_0 \cdot p_2^0$, $p^1 = p_0^1 \cdot c \cdot p_1^1 \cdot r_1 \cdot p_2^1$, call action c is matched by r_0 in p^0 , and is matched by r_1 in p^1 . Let $p_1^0 = a_0, b_0^0, a_1, b_1^0, \dots, a_n, b_n^0$, and let $p_1^1 = a_0, b_0^1, a_1, b_1^1, \dots, a_n, b_n^1$. Then $b_i^0 = b_i^1$ for each i in each such p^0 and p^1 . \square

Defn. ?? formalizes the informal policy-weaving problem illustrated in §??. As discussed in §??, if a policy-weaving problem has a solution I , then it has a solution I^* that may be represented as a VPA transducer T (i.e., a VPA where each action is labeled with input and output symbols). The problem of rewriting program P to satisfy the policy thus amounts to applying T to P, using a standard product construction from automata theory.

Privilege-aware systems are typically applied to monitor programs that could run injected code, yet an instrumentation function is defined in Defn. ?? to choose a primitive after each instruction executed by the program. However, this is not a fundamental limitation, as if a programmer or static analysis tool

determines that injected code might be run at a particular point in the program, then we can define the monitor so that no primitive other than a `noop` can be invoked by the instrumentation. Conversely, it is not too restrictive to only allow an instrumentation function to invoke a single primitive after each instruction, as we can rewrite the program to execute a sequence of security-irrelevant instructions between which the instrumentation can invoke a sequence of primitives. In App. ??, we present two different privilege-aware systems as VPA.

3.2 From Policy Weaving to Safety Games

Each policy-weaving problem $\text{POLWEAVE}(\mathsf{P}, \mathsf{Pol}, \mathsf{M})$ can be reduced to a *single-entry VPA (SEVPA [?]) safety game* that accepts plays corresponding to instrumented executions of P that violate Pol when run on M . A SEVPA safety game is a VPA structured as a set of modules with unique entry points whose transitions are decided in turn by an Attacker and a Defender. The states of the game are partitioned into modules, where the system transitions to a unique module on each call transition.

Definition 3. A SEVPA *safety game* for Attacker internal actions $\Sigma_{I,A}$, Defender internal actions Σ_D , call actions Σ_C , and return actions Σ_R is a tuple $\mathsf{G} = (Q_A, Q_D, Q_0, \iota_0, \{(Q_c, \iota_c)\}_{c \in \Sigma_C}, Q_F, \tau_{I,A}, \tau_D, \tau_R)$, where

- $Q_A \subseteq Q$ is a finite set of *Attacker states*.
- $Q_D \subseteq Q$ is a finite set of *Defender states*. Q_A and Q_D partition the states of the game Q .
- Q_0 is the *initial module*.
- $\iota_0 \in Q_0 \cap Q_D$ is the *initial state*.
- For $c \in \Sigma_C$, Q_c is the *module of c* . The sets $\{Q_c\}_{c \in \Sigma_C}$ and Q_0 are pairwise disjoint, and partition Q .
- For $c \in \Sigma_C$, $\iota_c \in Q_c \cap Q_D$ is the *initial state of c* .
- $Q_F \subseteq Q_0 \cap Q_D$ is the set of *accepting states*.
- $\tau_{I,A} : Q_A \times \Sigma_{I,A} \rightarrow Q_D$ is the *Attacker internal transition function*.
- $\tau_D : Q_D \times \Sigma_D \rightarrow Q_A$ is the *Defender internal transition function*.
- $\tau_R : Q_A \times \Sigma_R \times (Q_A \times \Sigma_C) \rightarrow Q_D$ is the *return transition function*.

The modules are closed under internal transitions: for $x \in \{0\} \cup \Sigma_C$, $q \in Q_x$, and $a \in \Sigma_{I,A}$, $\tau_{I,A}(q, a) \in Q_x$, and for $a \in \Sigma_D$, $\tau_D(q, a) \in Q_x$. A SEVPA safety game is not defined by using an explicit call transition function, because each call on an action c pushes on the stack the calling Attacker state and calling action (we thus call $\Gamma = Q_A \times \Sigma_C$ the *stack symbols of the game*), and transitions to ι_c . The modules of a SEVPA safety game are closed under matching calls and returns: for $x \in \{0\} \cup \Sigma_C$, $c \in \Sigma_C$, $q_x \in Q_x$, $q_c \in Q_c$, and $r \in \Sigma_R$, $\tau_R(q_c, r, (q_x, c)) \in Q_x$.

The *plays* of a SEVPA are defined analogously to the traces of a VPA. Let the *configurations of G* be $C = Q \times \Gamma^*$, let the *attacker configurations* be $C_D = C \cap (Q_A \times \Gamma^*)$, and let the *defender configurations* be $C_D = C \cap (Q_D \times \Gamma^*)$. Let the *Attacker actions* be $\Sigma_A = \Sigma_{I,A} \cup \Sigma_C \cup \Sigma_R$. $\tau_A : C_A \times \Sigma_A \rightarrow C_D$ maps each Attacker configuration and Attacker action to a Defender configuration:

$$\begin{aligned} \tau_A((q, s), a) &= (\tau_{I,A}(q), s) && \text{for } a \in \Sigma_{I,A} \\ \tau_A((q, s), a) &= (\iota_c, (q, a) \cdot s) && \text{for } a \in \Sigma_C \\ \tau_A((q, s_0 \cdot s'), a) &= (\tau_R(q, a, s_0), s') && \text{for } a \in \Sigma_R \end{aligned}$$

Because each transition on a Defender action is to an Attacker state and each transition on an Attacker action is to a Defender state, all plays that transition to a defined configuration are in $(\Sigma_D \cdot \Sigma_A)^*$. Let $\rho : (\Sigma_D \cdot \Sigma_A)^* \rightarrow C_D$ map each play of alternating Defender and Attacker actions to the configuration that the game transitions to from reading the play: $\rho(\epsilon) = (\iota_0, \epsilon)$, and $\rho(p \cdot a \cdot b) = \tau_A(\tau_D(\rho(p), a), b)$. A play $p \in (\Sigma_D \cdot \Sigma_A)^*$ is accepted by \mathbf{G} if $\rho(p) = (q, \epsilon)$ with $q \in Q_F$. Let $\mathcal{L}(\mathbf{G})$ be the set of all plays accepted by \mathbf{G} . \square

Because all accepting states of a game are in the initial module, a game can only accept matched plays. Superscripts denote the VPA or SEVPA game to which various components belong; e.g., $Q^{\mathbf{G}}$ are the states of SEVPA game \mathbf{G} .

A *Defender strategy* of a two-player safety game \mathbf{G} is a function $\sigma : (\Sigma_A^{\mathbf{G}})^* \rightarrow \Sigma_D^{\mathbf{G}}$ that takes as input a sequence of Attacker actions, and outputs a Defender action. σ is a *winning strategy* if as long as the Defender uses it to choose his next transition of the game, the resulting play is not accepted by \mathbf{G} : formally, $\sigma_{\text{tr}}(\Sigma_A^{\mathbf{G}}) \cap \mathcal{L}(\mathbf{G}) = \emptyset$ (for σ_{tr} as defined in Defn. ??). Let σ be modular if it satisfies the condition analogous to a modular instrumentation function (Defn. ??).

Theorem 1. *For each policy-weaving problem $\mathcal{P} = \text{POLWEAVE}(\mathbf{P}, \text{Pol}, \mathbf{M})$, there is a SEVPA safety game $\mathcal{G} = \text{PolWeaveGame}(\mathbf{P}, \text{Pol}, \mathbf{M})$ such that each instrumentation function that satisfies \mathcal{P} defines a winning modular Defender strategy of \mathcal{G} , and each winning modular Defender strategy of \mathcal{G} defines a satisfying instrumentation function of \mathcal{P} .*

The intuition behind the construction of \mathcal{G} from $\mathcal{P} = \text{POLWEAVE}(\mathbf{P}, \text{Pol}, \mathbf{M})$ is given in §??. From \mathbf{P} , we construct a game $\mathbf{G}_{\mathbf{P}}$ that accepts all instrumented privileged instrumented executions of \mathbf{P} . From Pol , we constructed a game \mathbf{G}_{Pol} that accepts all instrumented privileged executions that violate Pol . We construct \mathcal{G} as the product of $\mathbf{G}_{\mathbf{P}}$, \mathbf{G}_{Pol} , and $\mathbf{G}_{\mathbf{M}}$. Proofs of all theorems stated in §?? and §?? are in App. ??.

4 Solving SEVPA Safety Games with Scaffolds

In this section, we present an algorithm `ScafAlgo` that finds a winning modular Defender strategy to a given SEVPA safety game. The algorithm uses an additional, potentially smaller game, which we call a *scaffold*. We present `ScafAlgo` as a non-deterministic algorithm, and demonstrate that a symbolic implementation builds a formula whose size is decided entirely by the size of the scaffold and an additional, tunable independent parameter. We describe a known algorithm for finding modular strategies [?] and a known symbolic algorithm for finding strategies of bounded size [?] as instances of `ScafAlgo`.

4.1 Definition and Key Properties of Scaffolds

The key characteristic of our algorithm is that it finds a winning Defender strategy to a given game using an additional game, called a scaffold, and a specified relation between the states of the scaffold and the states of the game.

Definition 4. (Scaffolds) Let \mathbf{S} and \mathbf{G} be two SEVPA safety games defined for Attacker actions $\Sigma_{I,A}$, Defender actions Σ_D , call actions Σ_C , and return actions Σ_R . \mathbf{S} is a *scaffold* of \mathbf{G} under $\mathcal{R} \subseteq Q^{\mathbf{S}} \times Q^{\mathbf{G}}$ if and only if:

1. If $q_S \in Q_F^S$ and for $q_G \in Q^G$, $\mathcal{R}(q_S, q_G)$, then $q_G \in Q_F^G$.
2. For $c \in \Sigma_C$, $\mathcal{R}(l_c^S, l_c^G)$.
3. For $a \in \Sigma_{I,A}$, $q_S \in Q_A^S$, and $q_G \in Q_A^G$, if $\mathcal{R}(q_S, q_G)$, then $\mathcal{R}(\tau_{I,A}(q_S, a), \tau_{I,A}(q_G, a))$.
4. For $a \in \Sigma_D$, $q_S \in Q_D^S$, and $q_G \in Q_D^G$, if $\mathcal{R}(q_S, q_G)$, then $\mathcal{R}(\tau_D(q_S, a), \tau_D(q_G, a))$.
5. For $c \in \Sigma_C$, $q_c^S \in Q_C^S$, $q_c^G \in Q_C^G$, $q^S \in Q^S$, $q^G \in Q^G$, if $\mathcal{R}(q_c^S, q_c^G)$ and $\mathcal{R}(q^S, q^G)$, then $\mathcal{R}(\tau_R(q^S, r, (q_c^S, c)), \tau_R(q^G, r, (q_c^G, c)))$. \square

If so, then \mathcal{R} is a *scaffold relation* from S to G .

Each scaffolding relation \mathcal{R} defines an Attacker simulation [?], but there are SEVPA related by an Attacker simulation that are not related by any scaffold relation. However, scaffold relations and modular strategies are connected by the following key property, which provides the foundation for our algorithm. First, we define an (S, \mathcal{R}, k) -strategy of a game G , which intuitively is a strategy whose structure tightly corresponds to a scaffold S , according to a relation \mathcal{R} from the states of S to those of G . For a game G and $Q' \subseteq Q^G$ such that $\{l_c\}_{c \in \Sigma_C} \subseteq Q' \subseteq Q^G$, let the *subgame* $G|_{Q'}$ be the game constructed from restricting the states and transition functions of G to the states in Q' . Each subgame G' of G defines a strategy $\sigma_{G'}$ as a VPA transducer. To compute $\sigma_{G'}(a_0, a_1, \dots, a_n)$, $\sigma_{G'}$ uses a_0, a_1, \dots, a_n as the Attacker actions for a play of G' . If G' is in an attacker state, then σ_G transitions it on the next unread a_i , and if G' is in a Defender state, then $\sigma_{G'}$ picks a fixed Defender transition of G' on which to transition. $\sigma_{G'}$ outputs the Defender action chosen by G' after it reads all of a_0, a_1, \dots, a_n . This construction is standard, and we leave a full definition for App. ??.

Definition 5. For sets A and B , let a relation $\mathcal{R} \subseteq A \times B$ be *k-image-bounded* if for each $a \in A$, $|\{b \mid b \in B, \mathcal{R}(a, b)\}| \leq k$. Let G be a game, let S be a scaffold of G under $\mathcal{R} \subseteq Q^S \times Q^G$, and let $k \in \mathbb{Z}$. An (S, \mathcal{R}, k) -Defender strategy σ' of G is a Defender strategy such that for some $\mathcal{R}' \subseteq \mathcal{R}$, $\mathcal{R}' \cap (Q_A^S \times Q_A^G)$ is *k-image-bounded*, $G' = G|_{\text{Im}(\mathcal{R}')}$, and $\sigma' = \sigma_{G'}$. \square

Let game G have a winning Defender strategy, and let S be a scaffold of G under a scaffold relation \mathcal{R} . Then S is a scaffold of some subgame of G' that defines a winning strategy of G , under a finer scaffold relation than \mathcal{R} .

Theorem 2. *Let G have a winning modular Defender strategy, and let S be a scaffold of G under $\mathcal{R} \subseteq Q^S \times Q^G$. Then for some k , there is a winning modular (S, \mathcal{R}, k) -Defender strategy of G .*

4.2 An Algorithm Parametrized on Scaffolds

To find modular winning Defender strategies to games, we can apply Thm. ?? to search for (S, \mathcal{R}, k) -strategies. The algorithm `ScafAlgo`, given in Alg. ??, takes a game G , scaffold S , relation $\mathcal{R} \subseteq Q^S \times Q^G$, and parameter k , and searches for an (S, \mathcal{R}, k) -strategy by searching for an $\mathcal{R}' \subseteq \mathcal{R}$ that satisfies the condition given in Defn. ??.

`ScafAlgo` searches for such an \mathcal{R}' in three main steps. In the first step, `ScafAlgo` non-deterministically chooses a *k-image-bounded* subrelation of \mathcal{R} from the Attacker states of S to the Attacker states of G . Specifically, on line [??], `ScafAlgo`

Input: G : a VPA safety game.
 S : a scaffold of G
 $\mathcal{R} \subseteq Q^S \times Q^G$: a scaffold relation.
Output: If G has a winning (S, \mathcal{R}, k) -strategy, then it returns a winning (S, \mathcal{R}, k) -strategy. Otherwise, it returns \perp .

```

/* Choose  $\mathcal{R}_A$ : a  $k$ -image-bounded subrelation of  $\mathcal{R}$  that defines
Attacker states of a candidate strategy. */
1  $\mathcal{R}_A := \text{nd-bounded-subrel}(\mathcal{R} \cap (Q_A^S \times Q_A^G), k)$ ;
/* Construct  $\mathcal{R}_D$ : a relation to Defender states of the candidate
strategy defined by  $\mathcal{R}_A$ . */
2  $\mathcal{R}_{D,\iota} := \{(v_c^S, t_c^G) \mid c \in \Sigma_C^G\}$ ;
3  $\mathcal{R}_{D,i} := \{(\tau_{I,A}^S(p_A, a), \tau_{I,A}^G(q_A, a)) \mid (p_A, q_A) \in \mathcal{R}_A, a \in \Sigma_{I,A}^G\}$ ;
4  $\mathcal{R}_{D,r} := \{(\tau_R^S(p_A, a, s_A), \tau_R^G(q_A, a, t_A)) \mid (p_A, q_A), (s_A, t_A) \in \mathcal{R}_A, a \in \Sigma_R^G\}$ ;
5  $\mathcal{R}_D := \mathcal{R}_{D,\iota} \cup \mathcal{R}_{D,i} \cup \mathcal{R}_{D,r}$ ;
/* Check if the candidate strategy defined by  $\mathcal{R}_A$  and  $\mathcal{R}_D$  is a
winning strategy. */
6  $\text{StrWins} := \forall (p_D, q_D) \in \mathcal{R}_D : q_D \notin Q_F \wedge \exists a \in \Sigma_D : (\tau_D^S(p_D, a), \tau_D^G(q_D, a)) \in \mathcal{R}_A$ ;
7 if StrWins then return  $\sigma_{G|_{\text{Img}(\mathcal{R}_A \cup \mathcal{R}_D)}}$  else return  $\perp$ 

```

Algorithm 1. ScafAlgo: non-deterministic algorithm that takes a game G , scaffold S , and relation $\mathcal{R} \subseteq Q^S \times Q^G$, and finds a winning modular Defender (S, \mathcal{R}, k) -strategy of G .

defines such a relation $\mathcal{R}_A \subseteq Q_A^S \times Q_A^G$ by calling a function $\text{nd-bounded-subrel} : (Q_A^S \times Q_A^G) \times \mathbb{Z} \rightarrow (Q_A^S \times Q_A^G)$, where $\text{nd-bounded-subrel}(\mathcal{R} \cap (Q_A^S \times Q_A^G), k)$ is a k -image-bounded subrelation of $Q_A^S \times Q_A^G$.

In the second step (lines [??]–[??]), ScafAlgo constructs a relation $\mathcal{R}_D \subseteq Q_D^S \times Q_D^G$ such that if there is any $\mathcal{R}^* \subseteq Q_D^S \times Q_D^G$ such that $G|_{\text{Img}(\mathcal{R}_A \cup \mathcal{R}^*)}$ defines a winning strategy of G , then the *candidate strategy* defined by $G|_{\text{Img}(\mathcal{R}_A \cup \mathcal{R}_D)}$ is a winning strategy of G . On line [??], ScafAlgo defines $\mathcal{R}_{D,\iota} \subseteq Q_D^S \times Q_D^G$ that relates each module-initial state of S to its corresponding module-initial state in G . On line [??], ScafAlgo defines $\mathcal{R}_{D,i} \subseteq Q_D^S \times Q_D^G$ that, for each $(p_A, q_A) \in \mathcal{R}_A$ and internal Attacker action $a \in \Sigma_{I,A}^G$, relates the a -successor of p_A to the a -successor of q_A . On line [??], ScafAlgo defines $\mathcal{R}_{D,r} \subseteq Q_D^S \times Q_D^G$ that, for each $(p_A, q_A), (s_A, t_A) \in \mathcal{R}_A$ and return action $a \in \Sigma_R^G$, relates the r -successor of (p_A, s_A) to the r -successor of (q_A, t_A) . On line [??], ScafAlgo defines $\mathcal{R}_D \subseteq Q_D^S \times Q_D^G$ as the union of $\mathcal{R}_{D,\iota}$, $\mathcal{R}_{D,i}$, and $\mathcal{R}_{D,r}$.

In the third step (lines [??] and [??]), ScafAlgo checks if the candidate strategy defined by $G|_{\text{Img}(\mathcal{R}_A \cup \mathcal{R}_D)}$ is a winning strategy of G . On line [??], ScafAlgo defines $\text{StrWins} : \mathbb{B}$, which is true if and only if for each Defender-state of the candidate strategy, the state is not an accepting state of the game, and there is some action that the Defender can take to reach some Attacker-state of the candidate strategy. On line [??], ScafAlgo returns the strategy defined by $G|_{\text{Img}(\mathcal{R}_A \cup \mathcal{R}_D)}$ if and only if $G|_{\text{Img}(\mathcal{R}_A \cup \mathcal{R}_D)}$ is a winning strategy. Otherwise, ScafAlgo returns failure.

Theorem 3. Let G be a game, let S be a scaffold of G under $\mathcal{R} \subseteq Q^S \times Q^G$, and let k be a positive integer. If $\sigma = \text{ScafAlgo}(G, S, \mathcal{R}, k)$, then σ is a winning

Defender strategy for G . If G has a winning Defender strategy, then for each scaffold S and scaffolding relation $\mathcal{R} \subseteq Q^S \times Q^G$, there is some k such that $\text{ScafAlgo}(G, S, \mathcal{R}, k)$ is a winning Defender strategy of G .

A deterministic implementation of `ScafAlgo` runs in worst-case time exponential in the number of Attacker states. However, a symbolic implementation of `ScafAlgo` can represent its input problem with a formula whose size depends only on the scaffold, and the tunable parameter k . Assume that each component of an input game G is given as interpreted symbolic functions and predicates (i.e., states and actions are given as domains, and the transition functions are given as interpreted functions), and that the relation \mathcal{R} is given as an interpreted relation. Then `ScafAlgo` can be implemented by reinterpreting its steps to build a symbolic formula `StrWins` (line [??]) whose models correspond to values of \mathcal{R}_A and \mathcal{R}_D for which $G|_{\text{img}(\mathcal{R}_A \cup \mathcal{R}_D)}$ defines a winning strategy.

The size (i.e., the number of literals in) of `StrWins` is determined by S and k . The universal quantification on line [??] is bounded, and can thus be encoded as a finite conjunction; the nested existential quantification can then be Skolemized. To check the membership $(\tau_D^S(p_D, a), \tau_D^G(q_D, a)) \in \mathcal{R}_A$, we can apply the fact that \mathcal{R}_A is a k -bounded-image relation to represent the membership check with k disjuncts. From these observations, the size of the `StrWins` formula can be bounded by $O(|Q_A^S|^2 k^3)$.

Two known algorithms for finding modular strategies can be defined as `ScafAlgo` applied to degenerate scaffolds. A naive implementation of the original algorithm presented for finding modular strategies [?] can be defined as `ScafAlgo` applied to the game itself as a scaffold. A symbolic algorithm for finding strategies of bounded size [?], generalized to VPA games, can be defined as `ScafAlgo` applied to a scaffold with a single Attacker and Defender state for each module. The known algorithms are thus `ScafAlgo` applied to scaffolds that have complete and no information about their games, respectively. However, any game defined as a product of “factor” games, such as games constructed from the policy-weaving reduction, is scaffolded by its factors under a relation that relates each factor state to the product states for which it is a component. See App. ?? for a further discussion of known algorithms as instances of `ScafAlgo`.

5 Experiments

In this section, we discuss experiments that evaluate the reduction from policy-weaving problems to safety games presented in §??, and the scaffold-based game-solving algorithm presented in §??. The experiments were designed to answer two questions. First, by reducing policy-weaving problems to solving games, can we efficiently instrument practical programs for a real privilege-aware system so that they satisfy practical high-level policies? Second, which scaffolds allow our scaffolding game-solving algorithm to most efficiently solve games constructed by our policy-weaving algorithm?

To answer these questions, we instantiated our policy-weaving algorithm to a policy weaver for the Capsicum [?] capability operating system. We collected a set of six UNIX utilities, given in Tab. ??, that have exhibited critical security

Name	LoC	Pol.	Scaffolds							
			States		Triv.		Prog.-Pol.		Prog.-Pol.-Mon.	
			k	Time	k	Time	k	Time	k	Time
bzip2-1.0.6	8,399	12	12	-	1	0:04	1	0:09		
fetchmail-6.3.19	49,370	12	7	-	1	1:13	1	1:39		
gzip-1.2.4	9,076	9	12	-	1	1:47	1	-		
tar-1.25	108,723	12	3	3:47	1	1:20	1	-		
tcpdump-4.1.1	87,593	12	15	-	1	0:30	1	0:45		
wget-1.12	64,443	21	7	0:43	1	0:25	1	18:59		

Table 1. Performance of the Capsicum policy weaver. Column “LoC” contains lines of C source code (not including blank lines and comments), “Pol. States” contains the number of states in the policy, k contains the simulation bound, and “Times” contains the times used to find a strategy using the complete, trivial, and factor scaffolds. “-” denotes a time-out of 20 minutes.

vulnerabilities [?, ?, ?, ?]. For each utility, we defined a policy that describes the capabilities that the program must have as it executes. The policies were defined by working with the Capsicum developers, or using general knowledge of the utility. Detailed descriptions of the policies for each utility are given in [?].

We applied our Capsicum policy-weaver to each utility and its policy, each scaffold defined as a product of some subset of the program, policy, and monitor. The data from all scaffolds is given in App. ??; Tab. ?? presents data for several illustrative scaffolds: the trivial scaffold “Triv.” defined in §??, the product of the program and policy “Prog.-Pol”, and the product of all program, policy, and monitor “Prog.-Pol.-Mon.” For each scaffold, we measured how long it took our weaver to find a strategy, and with what minimum simulation bound (i.e., value of k from §??) it either found a strategy or timed out. The results for each scaffold are in the subcolumns of “Scaffolds” in Tab. ??, with each simulation bound in subcolumn “ k ,” and each time in subcolumn “Time.”

The results indicate that while many scaffolds give similar results for some practical problems, an *intermediate* scaffold constructed as a product of some but not all of the inputs, e.g. Prog.-Pol., leads to the best performance. The difference in performance could be due to fact that a scaffold with little information about the structure of its game (e.g., “Triv.”) generates a formula that allows many transitions between a small set of states in a candidate strategy, while a scaffold with total information (e.g., Prog.-Pol.-Mon.) generates a formula that allows few transitions between a large set of states in a candidate strategy. An intermediate scaffold strikes a balance between the two, generating a formula that allows a moderate number of transitions between a moderate set of states. We further discuss the experimental results in App. ??.

6 Related Work

Privilege-aware operating systems: Decentralized Information Flow Control (DIFC) operating systems such as Asbestos [?], HiStar [?], and Flume [?] manage privileges describing how information may flow through a system, and provide primitives that allow an application to direct flows by managing the labels of each

object in the system. Tagged memory systems such as Wedge [?] enforce similar policies per byte of memory by providing primitives for managing memory tags. Capability operating systems such as Capsicum [?] track the capabilities of each process, where a capability is a file descriptor paired with an access right, and provide primitives that allow an application to manage its capabilities.

Our work complements privilege-aware operating systems by allowing a programmer to give an explicit, high-level policy, and automatically rewriting the program to satisfy the policy when run on the system. Prior work in aiding programming for systems with security primitives automatically verifies that a program instrumented to use the Flume primitives enforces a high-level policy [?], automatically instruments programs to use the primitives of the HiStar to satisfy a policy [?], and automatically instruments programs [?] to use the primitives of the Flume OS. However, the languages of policies used in the approaches presented in [?,?] are not temporal and cannot clearly be applied to other systems with security primitives, and the proofs of the correctness of the instrumentation algorithms are ad hoc. The work in [?] describes the approach in this paper instantiated to a policy weaving for Capsicum. This paper describes how the work in [?] may be generalized to arbitrary privilege aware systems, and describes the novel game-solving algorithm applied in [?].

Inlined Reference Monitors: An *Inlined Reference Monitor (IRM)* [?,?] is code that executes in the same memory space as a program, observes the security-sensitive events of the program, and halts the program immediately before it violates a policy. IRMs have shortcomings that prohibit them from monitoring many practical programs and policies. Because an IRM executes in the same process space as the program it monitors, it cannot enforce policies throughout the system. Furthermore, an IRM must be able to monitor security-sensitive events of a program throughout the program’s execution, but there are known techniques to subvert an IRM [?]. Privilege-aware operating systems address the shortcomings of IRM by monitoring policies in the operating system, and providing a set of primitives that an application invokes to direct the operating system. The primitives are distinct from the security-sensitive events of interest.

Safety Games: Automata-theoretic games formalize problems in synthesizing reactive programs and control mechanisms [?]. Alur et. al. give an algorithm that takes a single-entry recursive state machine and searches for a strategy that is modular, as defined in §??, and show that this problem is NP-complete [?]. Recursive state machines are directly analogous to SEVPA [?]. Madusudan et. al. give a set of symbolic algorithms that find a winning strategy for a given game whose transition relation is represented symbolically [?]. The practical contribution of our work is that we express the emerging, important, and practical problem of rewriting programs for privilege-aware operating systems in terms of such games. We also give an algorithm for finding modular strategies that can be instantiated to a symbolic implementation of the algorithm of [?], to the “bounded-witness” algorithm of [?].

A Practical Privilege-Aware Systems as Automata

Practical privilege-aware systems can often be modeled as privilege-aware monitors as presented in §???. In this section, we informally model two apparently different privilege-aware systems, the Capsicum capability system [?] and the HiStar decentralized information-flow control (DIFC) system [?].

A.1 Capsicum as a VPA

The Capsicum capability system is the inspiration for the MiniCap system presented in §???, but is a robust, practical system included in the “RELEASE” version of FreeBSD 9.0 [?]. The full implementation of Capsicum extends the standard UNIX file descriptor to a *capability*, which is a descriptor paired with a set of *access rights*, which denote the ability of a process to read, write, seek, etc. a file. Capsicum defines 63 rights, significantly enriching the standard set of rights provided by UNIX. As each process executes, Capsicum maps the process to its set of capabilities, and only allows a process to access a file in a particular way (e.g., to read it) if the process has a sufficient set of capabilities. When a process begins to execute, it may open descriptors to any files available to the user, allowing the process to grant itself coarse sets of capabilities. As a process executes, it may invoke a primitive that restricts its capabilities for a file descriptor, and it may invoke another primitive to enter “capability mode”, which causes Capsicum to disallow it from opening additional descriptors. The primitives were designed for a programmer to rewrite their program to execute trusted code with high privilege, then restrict its capabilities by invoking the primitives, and then execute untrusted code with capabilities that are sufficiently reduced such that even a malicious or compromised module invoked by the program cannot violate a particular security policy.

Capsicum may be modeled as a privilege-aware system monitor as defined in §???. The privileges of the monitor are sets of capabilities. The primitives of the monitor are the system calls that a program invokes to restrict its capabilities or enter capability mode. The state of the monitor is the set of capabilities held by the program, and a Boolean that decides whether the program is in capability mode. A full description of Capsicum as a VPA monitor is given in [?].

A.2 HiStar as a VPA

The HiStar DIFC system [?] maps every process to a *label*, interposes on each time that one process attempts to send data to a receiving process, and allows the data to be sent if and only if the labels of the sending and receiving process satisfy a particular condition. A label is a finite set of named *categories*, which each may have a value in the range 0 – 3, or the special value \star . Intuitively, each category corresponds to a source of information, and a process with a higher numeric value for a category has read more sensitive information from the source. If a process has the \star value for a category c , then it has the ability to declassify information corresponding c . The names of the categories in the label of each process are the same, and HiStar allows a process p to send information to process q if and only if for each category c , the value of c in the label of p is

a numeric value less than or equal to the numeric value of c in p , or the value of p or q for c is \star .

HiStar provides system calls that each process may invoke to manipulate labels, subject to particular rules. A process may invoke one primitive to introduce a new category c , giving itself \star for c and assigning a default value of c for every other process. A process may also invoke a primitive to raise the numeric value of its label for any category.

HiStar may be modeled as a privilege-aware system monitor as defined in §???. The privileges of HiStar are the flows that it allows, and the primitives are the system calls that allow each process to manipulate its label. Each state of the HiStar VPA is a map from each process to its label.

B Proofs of Theorems

In this section, we sketch informal proofs of the theorems stated in §?? and §??.

Proof. (Thm. ??) First, assume that the languages accepted by the program automaton P , policy Pol , and monitor M only contain matched words. Then from each automaton, we construct a (non-game) SEVPA that accepts the language of the automaton [?]. We then construct a pair of games whose Attacker actions are program instructions, and whose Defender actions are monitor primitives. Let G_P be a SEVPA game that accepts all instrumented executions of P : $\mathcal{L}(G_P) = \{[p_0, i_0, \dots, p_n, i_n] \mid [i_0, \dots, i_n] \in \mathcal{L}(P), p_j \in \text{Prims}\}$. Let $G_{Pol, M}$ be a SEVPA game that accepts all instrumented executions for M that violate Pol : $\mathcal{L}(G_{Pol, M}) = \{[p_0, i_0, \dots, p_n, i_n] \mid \exists r_j \in \text{Privs} : [p_0, (i_0, r_0), \dots, p_n, (i_n, r_n)] \in \mathcal{L}(M) \wedge [(i_0, r_0), \dots, (i_n, r_n)] \in \mathcal{L}(Pol)\}$. Let G be such that $\mathcal{L}(G) = \mathcal{L}(G_P) \cap \mathcal{L}(G_{Pol, M})$, and let $\text{PolWeaveGame}(P, Pol, M) = G$. It is straightforward to show that the instrumentation functions of $\text{POLWEAVE}(P, Pol, M)$ are exactly the modular winning Defender strategies of $\text{PolWeaveGame}(P, Pol, M)$.

The direct construction of $G_{Pol, M}$ may allow $G_{Pol, M}$ to be non-deterministic. To avoid this in practice, we require that the M in fact be a *transducer function* from instrumented executions to privileges. Using this assumption, one can directly construct a deterministic $G_{Pol, M}$ from deterministic Pol and M without requiring an explicit determinization.

To prove Thm. ??, we first recall a well-known construction of strategies, stated informally in §??. For a SEVPA safety game G with subgame G' , define the strategy $\sigma_{G'}$ as follows. For each Defender state $q \in Q_D^{G'}$, fix a Defender action $\delta(q)$ such that $\tau_D(q, \delta(q)) \in Q^{G'}$. Let $\gamma_{G'} : (\Sigma_A^G)^* \rightarrow C^{G'}$ map each trace of Attacker actions to the configuration of G' that the actions interleaved with fixed Defender actions drive G' : $\gamma_{G'}(\epsilon) = (\iota_0, \epsilon)$, and $\gamma_{G'}(s \cdot a) = \tau_A(\tau_D(\gamma_{G'}(s), \delta(\text{St}(\gamma_{G'}(s))), a)$, where $\text{St}(c)$ is the state of configuration c . Then $\sigma_{G'}(s) = \delta(\text{St}(\gamma_{G'}(s)))$.

If a SEVPA game has a modular winning Defender strategy, then it has a modular winning strategy $\sigma_{G'}$ for some subgame G' of G . This follows immediately from results for finding modular strategies to recursive game graphs [?].

Proof. (Thm. ??) Let G have a modular winning Defender strategy σ_{G^*} , where G^* is a subgame of G . Let $R' = R \cap (Q^S \times Q^{G^*})$. For $G' = G|_{\text{img}(R')}$, $\sigma_{G'}$ is a

modular winning Defender strategy, by a straightforward argument. Thus for $k = \max_{q \in Q_A^S} |\{q' | R'(q, q')\}|$, G' is an (S, R, k) strategy, where R' satisfies the condition for R' given in Defn. ??.

Proof. (Thm. ??) *Soundness:* Let $G' = G|_{\text{Img}(\mathcal{R}_A \cup \mathcal{R}_D)}$, for \mathcal{R}_A and \mathcal{R}_D as defined in Alg. ??, it is straightforward to show that $\sigma_{G'}$ is a winning strategy using the definition of **StrWins**. The modularity of $\sigma_{G'}$ follows from the fact that it is a subgame of a SEVPA.

Completeness: If G has a winning strategy, then by Thm. ??, it has an (S, \mathcal{R}, k) winning modular strategy for some k . Let this (S, \mathcal{R}, k) be σ_{G^*} for a $G^* = G|_{\text{Img}(\mathcal{R}')}$, where $\mathcal{R}' \subseteq \mathcal{R}$, and let $\mathcal{R}_A = \mathcal{R}' \cap (Q_A^S \times Q_A^G)$. Let **ScafAlgo** non-deterministically guess \mathcal{R}_A on Alg. ?? line [??], and then construct \mathcal{R}_D on line [??]. It follows that $\text{Img}(\mathcal{R}_D) \subseteq Q^{G'}$, and thus that $\text{Img}(\mathcal{R}_D) \cap Q_F^G = \emptyset$ and each state in $\text{Img}(\mathcal{R}_D)$ transitions on some Defender action to some state in $\text{Img}(\mathcal{R}_A)$. It then follows that for $G' = G|_{\text{Img}(\mathcal{R}_A) \cup \text{Img}(\mathcal{R}_D)}$, $\sigma_{G'}$ is a winning modular strategy of G .

C Known Game Algorithms as Scaffolds

Various algorithms have been presented for finding modular winning strategies to games. We now show that a naive version of a known algorithm for finding modular strategies for recursive game graphs [?] and a known symbolic algorithm for finding strategies of bounded size for symbolic games [?] are analogous to instances of **ScafAlgo**.

Symbolic Implementation of the Standard Algorithm Alur et. al. gave an explicit algorithm for finding modular Attacker strategies to recursive game graphs [?], and proved that the problem of finding modular strategies is NP-complete. Thus the algorithm may run in time exponential in the size of the game. The algorithm given in [?] can be adopted to find modular Defender strategies to SEVPA games, and we call the adopted algorithm **Standard**. **Standard** finds, for each state q , a set of states X_q such that some modular strategy can ensure that if the game reaches q , then it always stays within X_q , and X_q does not contain an accepting state. **Standard** then picks a set of Defender transitions that ensure that the game stays in X_q .

To compute X_q for each state s , **Standard** computes a set of sets of candidate X_q , and as a result, **Standard** may run in time exponential in the size of an input game G . **ScafAlgo** can analogously compute the sets X_q using the input game itself as a scaffold. Let $=_{Q^G}$ be the equality relation over the states of G . If we apply **ScafAlgo**($G, G, =_{Q^G}, 1$), then **ScafAlgo** non-deterministically guesses whether each state in G is in a candidate winning strategy, and then checks if a winning strategy can be constructed from the guessed Attacker states. A symbolic implementation of **Standard** using **ScafAlgo** constructs formulas of size $O(|Q_A^G|^2)$.

The algorithm presented in [?] is defined over recursive game graphs, in which the internal states of each game module are disjoint from entry and exit states. The algorithm executes in time exponential only in the number of *exit* states, not

Performance data of scaffolds										
Name	LoC	Pol. States	Scaffolds							
			Triv.		Prog.		Mon.		Pol.	
			k	Time	k	Time	k	Time	k	Time
bzip2-1.0.6	8,399	12	12	-	4	1:26	15	-	7	17:22
fetchmail-6.3.19	49,370	12	7	-	1	0:16	2	-	2	1:54
gzip-1.2.4	9,076	9	12	-	1	-	1	-	8	-
tar-1.25	108,723	12	3	3:47	1	-	1	-	3	5:17
tcpdump-4.1.1	87,593	12	15	-	4	1:13	23	-	6	-
wget-1.12	64,443	21	7	0:43	1	0:15	1	-	2	0:42

Performance data (cont.)									
Name	Scaffolds								
	Prog.-Mon.		Prog.-Pol.		Mon.-Pol.		Prog.-Mon.-Pol.		
	k	Time	k	Time	k	Time	k	Time	
bzip2-1.0.6	15	-	1	0:04	8	-	1		0:09
fetchmail-6.3.19	1	1:13	1	1:13	1	-	1		1:39
gzip-1.2.4	1	-	1	1:47	1	-	1		-
tar-1.25	1	-	1	1:20	1	-	1		-
tcpdump-4.1.1	18	-	1	0:30	9	-	1		0:45
wget-1.12	1	-	1	0:25	1	-	1		18:59

Table 2. Performance of the Capsicum policy weaver. Column “LoC” contains lines of C source code (not including blank lines and comments), “Pol. States” contains the number of states in the policy, k contains the simulation bound, and “Times” contains the times used to find a strategy using the complete, trivial, and factor scaffolds. “-” denotes a time-out of 20 minutes.

the number of all states. In this work, we use SEVPA’s to define games, which are analogous to recursive game graphs [?], but do not have disjoint sets of internal, entry, and exit states, and thus are not amenable to deriving such bounds.

Symbolic Search for Bounded Witness Sets The algorithm of [?] computes information for each state in a given game. However, many games have state spaces too large to enumerate explicitly. Madhusudan et. al. [?] describe how to construct, from a game G and bound k on the number of states that may be in a strategy, a constraint $\beta(G, k)$ whose models correspond to winning strategies with at most k states. Each strategy is constructed from a bounded *witness set* of states of the game [?], and thus we call the algorithm **BoundedWS**. The states in a witness set are derived from a model of $\beta(G, k)$.

We can construct $\beta(G, k)$ by applying **ScafAlgo**. For a game G , define a “trivial” scaffold S_T with a module for each call in Σ_C^G that contains a single Defender and Attacker state. Let \mathcal{R} relate the Defender state of each module Q_c^S for $c \in \Sigma_C^G$ of S_T to each Defender state in module Q_c^G , and relate the Attacker state of module Q_c^S to each Attacker state in module Q_c^G . Then the **StrWins** formula constructed by a symbolic implementation of **ScafAlgo** for **ScafAlgo**(G, S, \mathcal{R}, k) has models that directly correspond to the models of $\beta(G, k)$. The size of the formula constructed by **ScafAlgo**(G, S, \mathcal{R}, k) is $O(k^3)$.

D Complete Experimental Data

In §??, we described a set of experiments in which we implemented a policy weaver for the Capsicum operating system, and applied it to a set of UNIX utilities and the policies, using a variety of scaffolds to search for a winning strategy. In §??, we presented results for the two extreme scaffolds Triv. and Prog.-Pol.-Mon., and the intermediate scaffold Prog.-Pol. The tables in Table 2 present the results of applying all scaffolds constructed as subsets of products of the program, policy, and monitor automata. Each scaffold is labeled by the automata of which it is a product.

While many scaffolds provide performance comparable to the Prog.-Pol. scaffold on particular benchmarks, Prog.-Pol consistently finds strategies in minutes, whereas each other scaffold times out on at least two benchmarks. Prog.-Pol.’s consistency likely results from the fact that it always finds a winning strategy with a simulation bound of 1, whereas if `ScafAlgo` is applied to the other scaffolds, it often must use a higher simulation bound. Thus, apparently all of the practical policy-weaving problems that we found can be solved by assigning a single monitor state to each program-policy state pair, such that the strategy ensures that when an execution reaches the program and policy states, it will be in the assigned monitor state. In comparison, the scaffold Prog. sometimes requires a simulation bound of 4, indicating that some policy problems require finding up to 4 monitor-policy states for each program state.