

On the automated synthesis of proof-carrying temporal reference monitors

Simon Winwood^{1,2}, Gerwin Klein^{1,2}, and Manuel M. T. Chakravarty¹

¹ University of New South Wales
School of Computer Science & Engineering
Sydney, Australia

² National ICT Australia*

{sjw,chak}@cse.unsw.edu.au gerwin.klein@nicta.com.au

Abstract. We extend the range of security policies that can be guaranteed with proof carrying code from the classical type safety, control safety, memory safety, and space/time guarantees to more general security policies, such as general resource and access control. We do so by means of (1) a specification logic for security policies, which is the past-time fragment of LTL, and (2) a synthesis algorithm generating reference monitor code and accompanying proof objects from formulae of the specification logic. To evaluate the feasibility of our approach, we developed a prototype implementation producing proofs in Isabelle/HOL.

1 Introduction

Proof carrying code (PCC) [15] is inherently trustworthy, independent of its origin or previous opportunities for tampering. The guarantees provided by PCC are, however, not universal: they are relative to a *security policy* agreed upon by the code producer and consumer. It is the code producer’s obligation to annotate the code with a *proof object* that establishes the code’s compliance with the security policy. This proof object, consisting of steps in a *formal logic*, can be checked with a simple proof checker. Thus, the trustworthiness of the code can be established with mathematical rigour.

Existing research into the generation of proof-carrying code focuses on security policies which can be derived from properties of high-level languages and their type systems, such as type safety [14], control and memory safety [16], and space/time guarantees [1]. The contribution of this paper is to extend the approach to more general security policies, such as general resource and access control. An example of such a policy is one where “a user may perform an operation only if they have been granted a capability for that operation and that capability hasn’t been revoked.” Such properties are beyond the semantic guarantees of high-level languages; hence, we need (1) a formal device to express such policies and (2) a method for generating proof-carrying code for these policies.

* National ICT Australia is funded through the Australian Government’s *Backing Australia’s Ability* initiative, in part through the Australian Research Council

To address Point (1), we introduce a fragment of LTL [13] lacking the usual future operators (*until* and *next*), which we call *propositional pure-past temporal logic* (P3TL), as a specification logic for security policies in Section 3. P3TL can express a wide range of security policies, while enabling the automatic synthesis of reference monitors [20] that enforce P3TL policies. This second property is key to solving Point (2). More precisely, in Section 5, we will give an algorithm to synthesise, firstly, a reference monitor checking a given P3TL policy, and secondly, a proof object demonstrating that the reference monitor code indeed meets the policy. Such a reference monitor in conjunction with a framework to constrain application code to abide by the rules of a reference monitor is sufficient to produce PCC for P3TL policies. To be complete, this framework must also provide machine checkable proof that the application code cannot subvert the reference monitor. We introduced one such framework based on hybrid sandboxing accompanied by a proof in the theorem prover Isabelle/HOL in previous work [21]. We will describe this set up in more detail in Section 2.

The proof of compliance of synthesised reference monitors is in a Hoare-like logic discussed in Section 4 and formalised in Isabelle/HOL. We implemented a prototype synthesis tool in Isabelle/HOL to demonstrate the practical feasibility of our approach.

In summary, our specific contributions are these:

- a formalisation of a simple language and program logic for reference monitors (Section 4); and,
- a synthesis algorithm for reference monitors and proof objects from P3TL formulae (Section 5).

In contrast to previous work on generating reference monitors from temporal logics, we simultaneously generate a proof object that demonstrates that the generated code enforces the security policy. We discuss related work further in Section 6.

2 Our approach

Fig. 1 shows an application scenario in which our technique is useful. Assume a code producer generating a PCC program, certified for some policy ϕ . The code consumer, however, requires a stronger policy, ψ , where the additional guarantees of ψ over ϕ are beyond those that can be directly included during PCC generation. The method introduced in this paper enables the synthesis of a reference monitor for ψ ; this can then be inserted into the PCC program, possibly by rewriting parts of the application, to bridge the gap between ϕ and ψ . This paper presents an algorithm for synthesising such monitors with matching proofs. In previous work [21], we showed how such a monitor can be integrated into an existing application in the special case where that existing application does not use PCC at all; i.e., ϕ is empty. The general case, depicted in Fig. 1, where two policies need to be integrated is left for future work.

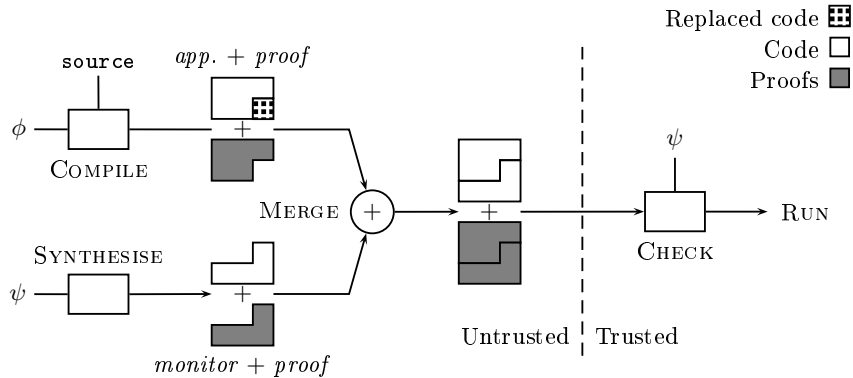


Fig. 1. A possible application of our approach: generate a monitor for the stronger policy ψ and insert it into the application. This paper presents the monitor and proof synthesis component.

2.1 The monitor environment

We require that the surrounding PCC infrastructure ensures the monitor invariant, containing assertions about the monitor state, be maintained outside the monitor; in a typical system, we expect the monitor state to be hidden from the rest of the system, either through a module mechanism or through general memory safety — the latter is the option we explored in [21]. If a to be secured operation occurs outside the monitor, the invariant will, in general, be invalidated — this will be detected at proof check time, and the program rejected. Of course, it is perfectly valid for there to be secured operations outside the monitor if the outside code can ensure that the invariant is maintained and that assurance is reflected in the code’s proof annotations.

2.2 The prototype

We have implemented a prototype synthesis tool. The tool generates an Isabelle theory file containing both the monitor code and proofs. The tool is Isabelle-specific, however we expect an implementation for another logical framework to have many similarities. The prototype, along with the Isabelle theory file, is freely available.³

While we make extensive use of Isabelle, we have been careful to ensure that there is no intrinsic reliance upon any particular feature. In the generated proofs, we avoid, where possible, use of Isabelle’s automated tactics, using them only when the lemma to be proved is Isabelle-specific. This includes, for example, lemmas relating to substitution: because we make use of the meta-logic’s substitution (we use a shallow embedding for the assertion logic), the proof will be Isabelle-specific, and thus no general proof is available. Conversely, the monitor’s proof obligations are independent of the particular meta-logic, and so we use individual proof rules for the main part of the proof.

³ <http://www.cse.unsw.edu.au/~sjw/papers/synthesis.html>

3 Policy Logic

In this section we introduce the logic we use for describing security policies.

Example 1. Consider a simple version of the Chinese Wall security policy[5]

A user may access files for any client, but once they have done so, they may only access files for that client.

Given an operator \diamond which denotes "at some time in the past", this policy may be expressed as⁴

$$\text{access}(f) \wedge f \in C \longrightarrow \neg \diamond (\text{access}(g) \wedge g \notin C)$$

which states a policy equivalent to the one above: access to an object belonging to a client is allowed only if at no point in the past has the user accessed an object belonging to another client.

Example 2. Now consider a policy modelling a capability system

A user may perform an operation only if they have been granted a capability for that operation and that capability hasn't been revoked.

We may encode this policy as follows

$$\text{operate}_o \longrightarrow \neg \text{revoke}_o \mathcal{S} \text{grant}_o$$

where $\psi \mathcal{S} \phi$ means ϕ was true at some point in the past, and ψ ever since.

These policies, like many others, express properties over a *series* of events: the first policy that some event had not occurred in the past; the second policy that some event had occurred, and in the meantime another had not.

Given our general goal of synthesising reference monitors, our policy logic needs to fulfill two requirements: (1) we must be able to use reference monitors to implement all policies that are expressible in the logic and (2) the logic must be powerful enough to express common policies, such as those above.

The latter means that the logic must be able to express properties over behaviours of the program, not just over individual states. The former means that the logic should express safety properties only, as liveness properties cannot be implemented by reference monitors [20].

These two requirements still leave some choice. We have settled on the safety fragment [13] of propositional linear temporal logic (LTL); for the remainder of the paper we will denote this fragment *propositional pure-past temporal logic* (P3TL). As the name signifies, this is a propositional logic containing temporal operators that refer solely to previous worlds. The logic differs from traditional temporal logics in that it lacks the usual future (*until* and *next*) operators — a

⁴ In this section, to simplify the formalisations, we are a little loose with syntax.

$$\begin{aligned}
\langle \sigma, n \rangle \models \langle \mathbf{a} \rangle &= \mathbf{a} \sigma_{[n]} \\
\langle \sigma, n \rangle \models \varphi [\Rightarrow] \psi &= \langle \sigma, n \rangle \models \varphi \longrightarrow \langle \sigma, n \rangle \models \psi \\
\langle \sigma, n \rangle \models \varphi \mathcal{S} \psi &= \exists i \leq n. \langle \sigma, i \rangle \models \psi \wedge (\forall j \in (i..n). \langle \sigma, j \rangle \models \varphi) \\
\langle \sigma, n \rangle \models \ominus \varphi &= n \neq 0 \longrightarrow \langle \sigma, n - 1 \rangle \models \varphi
\end{aligned}$$

Fig. 2. Semantics of P3TL

formula ψ in P3TL is equivalent to $\Box\psi$ in the LTL of Manna and Pnueli [13]. A similar logic (ptLTL) is used by Havelund and Rosu [9].

A more powerful logic, such as a first-order variant of P3TL, would enable more policies and more succinct versions of those already mentioned. Unfortunately, synthesising reference monitors for these logics is much more difficult than for P3TL — the state space for P3TL is fixed, while that of a first-order logic is potentially unbounded. We leave extension of our system to more powerful logics as future work.

The following is the syntax of P3TL

$$\text{form} ::= \langle \text{atom} \rangle \mid \text{form}_1 [\Rightarrow] \text{form}_2 \mid \text{form}_1 \mathcal{S} \text{form}_2 \mid \ominus \text{form}$$

The logic consists of atoms and the operators implication, since, and weak previous. Implication is the usual binary operator; to distinguish it from that of HOL, we write $\varphi [\Rightarrow] \psi$. The syntax for the since and weak previous operators is standard.

The definition of P3TL is parameterised over the type of states; atoms are HOL predicates on such states. Our implementation of P3TL is then a *deep embedding* into Isabelle/HOL, but uses a *shallow embedding* for atomic propositions.

Example 3. A traditional state space would be a tuple or record of variables. The predicate $\langle \lambda s. x \ s = 7 \rangle$ then, for instance, states that the variable x in record s should have value 7.

The following gives the definitions for the other propositional connectives (negation, conjunction, disjunction, etc.) as well as the usual temporal operators: *strong previous* ($\ominus \varphi$), *once* ($\diamond \varphi$), and *so-far* ($\square \varphi$). The two previous operators (strong and weak) differ only at the initial state: $\ominus \varphi$ is false and $\ominus \perp$ is true. In particular, $\ominus \perp$ is true only at the initial state.

$$\begin{array}{ll}
\top & = \langle \lambda s. \text{True} \rangle & \ominus \varphi & = [\neg] (\ominus ([\neg] \varphi)) \text{ (Previously)} \\
\perp & = \langle \lambda s. \text{False} \rangle & \diamond \varphi & = \top \mathcal{S} \varphi \text{ (Once)} \\
[\neg] \varphi & = \varphi [\Rightarrow] \perp & \square \varphi & = [\neg] (\diamond ([\neg] \varphi)) \text{ (So-far)} \\
\varphi [\wedge] \psi & = [\neg] (\varphi [\Rightarrow] [\neg] \psi) & \varphi [\vee] \psi & = [\neg] \varphi [\Rightarrow] \psi
\end{array}$$

Fig. 2 gives the semantics of P3TL using an *indexed model* $\langle \sigma, n \rangle$. The first element, σ , is a sequence of worlds, and the second, n , is the index of the current state.

An atom $\langle \mathbf{a} \rangle$ is valid in $\langle \sigma, \mathbf{n} \rangle$ iff the function \mathbf{a} maps the \mathbf{n} -th state in σ (denoted $\sigma_{[\mathbf{n}]}$) to true. An implication is valid iff the validity of the premise implies the validity of the conclusion at the same index \mathbf{n} . The formula $\varphi \mathcal{S} \psi$ is valid at $\langle \sigma, \mathbf{n} \rangle$ iff ψ was valid at some earlier state \mathbf{i} and φ was valid at all states \mathbf{j} from \mathbf{i} to \mathbf{n} . The weak previous operator $\odot \varphi$ is valid iff there is either no previous state ($\mathbf{n} = 0$) or φ was valid at index $\mathbf{n} - 1$.

4 A Language and Logic for Reference Monitors

4.1 The Programming Language

This section describes the language we use for reference monitors. It is a simple imperative if-while language, using Isabelle/HOL's functions for expressions. Whilst it is usual in expositions on proof-carrying code systems to use a very low-level language (e.g., an assembly language), we chose this comparatively high-level language to simplify the presentation; issues such as memory allocation are orthogonal to the contribution of this work and would only obscure the central points. We expect that low-level languages will present few additional theoretical hurdles — indeed, Barthe et. al. [3] note that, in the absence of optimisations, transformation to a low-level language preserves proof obligations.

The syntax of programs is shown below. The nonterminal *basic* denotes functions from states to states, *bexp* functions from states to booleans, and *form* P3TL formulae.

$$\begin{aligned} \text{stmt} ::= & \text{Do } \text{basic} \quad | \quad \text{stmt}_1 ; \text{stmt}_2 \\ & | \quad \text{IF } \text{bexp} \text{ THEN } \text{stmt}_1 \text{ ELSE } \text{stmt}_2 \text{ FI} \quad | \quad \text{WHILE } \text{bexp} \text{ DO } \text{stmt} \text{ OD} \\ & | \quad \text{Secure } \text{form} \end{aligned}$$

As with P3TL, the syntax is parameterised over the states of the program. The **Do** statement provides the means to model assignment and simple state transformations directly as HOL functions. The **Secure** statement represents the operation to be secured, abstracting away from the particular operation's semantics. It is parameterised by a P3TL formula representing the security policy.

Example 4. The following monitor checks the Chinese Wall policy from Example 1

```

IF ( $\lambda s. \exists f. \text{access } s \ f \wedge f \in \mathbb{C}$ ) THEN
  IF ( $\lambda s. \text{seen } s = 1$ ) THEN
    Do ( $\lambda s. s(\text{error} := 1)$ )
  ELSE
    Secure  $\psi$ 
  FI
ELSE
  Secure  $\psi$ 
FI

```

$$\boxed{s -t \rightarrow s'}$$

$$\frac{}{(s, ss) -\text{Do } f \rightarrow (f \ s, ss)}$$

$$\frac{s -e \rightarrow t \quad t -e' \rightarrow s'}{s -e; e' \rightarrow s'}$$

$$\frac{b \ (\text{fst } s) \quad s -e \rightarrow s'}{s -\text{IF } b \ \text{THEN } e \ \text{ELSE } e' \ \text{FI} \rightarrow s'}$$

$$\frac{\neg b \ (\text{fst } s) \quad s -e' \rightarrow s'}{s -\text{IF } b \ \text{THEN } e \ \text{ELSE } e' \ \text{FI} \rightarrow s'}$$

$$\frac{\neg b \ (\text{fst } s)}{s -\text{WHILE } b \ \text{DO } e \ \text{OD} \rightarrow s}$$

$$\frac{b \ (\text{fst } s) \quad s -e \rightarrow t \quad t -\text{WHILE } b \ \text{DO } e \ \text{OD} \rightarrow s'}{s -\text{WHILE } b \ \text{DO } e \ \text{OD} \rightarrow s'}$$

$$\frac{s' = (\text{if } (s, ss) \models \varphi \ \text{then } (s, ss @ [s]) \ \text{else arbitrary})}{(s, ss) -\text{Secure } \varphi \rightarrow s'}$$

Fig. 3. Semantics of programs.

The state variable `seen` records whether we have previously seen a conflicting access — the code to maintain this is omitted. The expression `s(error := 1)` updates `error` in record `s`.

The states over which the language operates are tuples (s, σ) , where s is the program state, usually modelled as an Isabelle record⁵, and σ a sequence representing a trace of previously seen states. This trace serves as the model when checking P3TL formulae; it records the state at security events only, not all state changes in the program. Note that the trace is only required for reasoning about policies and does *not* appear at runtime.

In monitoring P3TL policies, we interpret the validity of formulae relative to the current state, that is, the last state in the sequence. We thus use an *anchored interpretation* [12]. Satisfaction of P3TL formulae by a program state is then

$$(s, \sigma) \models \varphi = (\sigma @ [s], |\sigma|) \models \varphi$$

Fig. 3 shows a big step semantics for our language. The semantics are standard [17], apart from the `Secure` statement: the effect of `Secure` φ is to record the current state in the state history. Execution of `Secure` φ is only defined, however, when the current state satisfies the security policy φ .

A full implementation of our system would replace the `Secure` statement by, for example, a system call statement. Again, modelling the behaviour of system calls is orthogonal to the aims of this paper; the much-simplified `Secure` is sufficient.

⁵ As we use the `Do` statement to model assignment, taking a state update function as an argument, there is no requirement that a record is used — it merely simplifies use of the language.

$$\begin{array}{c}
\boxed{\Vdash \{P\} T \{Q\}} \\
\hline
\frac{}{\Vdash \{\lambda(s, ss). P (f s, ss)\} \text{Do } f \{P\}} \text{DO} \qquad \frac{\Vdash \{P\} e \{Q\} \quad \Vdash \{Q\} e' \{R\}}{\Vdash \{P\} e; e' \{R\}} \text{SEQ} \\
\hline
\frac{\Vdash \{\lambda s. P s \wedge b (fst s)\} e \{Q\} \quad \Vdash \{\lambda s. P s \wedge \neg b (fst s)\} e' \{Q\}}{\Vdash \{P\} \text{IF } b \text{ THEN } e \text{ ELSE } e' \text{ FI } \{Q\}} \text{COND} \\
\hline
\frac{\Vdash \{\lambda s. P s \wedge b (fst s)\} e \{P\} \quad Q = (\lambda s. P s \wedge \neg b (fst s))}{\Vdash \{P\} \text{WHILE } b \text{ DO } e \text{ OD } \{Q\}} \text{WHILE} \\
\hline
\frac{\forall s. P' s \longrightarrow P s \quad \forall s. Q s \longrightarrow Q' s \quad \Vdash \{P\} e \{Q\}}{\Vdash \{P'\} e \{Q'\}} \text{CONS} \\
\hline
\frac{\forall s. P s \longrightarrow s \models \varphi \quad \forall s ss. P (s, ss) \longrightarrow Q (s, ss @ [s])}{\Vdash \{P\} \text{Secure } \varphi \{Q\}} \text{SECURE}
\end{array}$$

Fig. 4. A Hoare-like logic for the programming language.

4.2 The Program Logic

So far we have defined a logic for policies and a language to implement reference monitors. This section introduces a logic for reasoning reasons about programs that allows us to formally verify that safety policies are respected.

The rules of this Hoare-like program logic are shown in Fig 4. The triple $\Vdash \{P\} T \{Q\}$ denotes that a statement T that starts execution in a state satisfying P and terminates will finish in a state satisfying Q . Both P and Q are assertions, that is, HOL predicates on states.

As with the semantics in Sect. 4.1, the proof rules are standard apart from the `Secure` statement. We have one rule for each syntactic construct and the usual rule of consequence.

The new rule for `Secure` demands that all states that satisfy the precondition be models of the security policy φ ; this check links the assertion logic to P3TL, the policy logic. In addition, the rule's postcondition reflects the effect of `Secure` on the program state, i.e. recording the event in the trace.

Following the standard practise [17] of using a shallow embedding of assertions into Isabelle/HOL means that we can take direct advantage of Isabelle's tactics and libraries to reason about programs. The assertion logic and program expressions, however, are more powerful than strictly required⁶.

We have shown this Hoare-logic to be sound in the following sense.

Theorem 1. If $\Vdash \{P\} e \{Q\}$ and $P s$ and $s \dashv\vdash e \rightarrow s'$ then $Q s'$.

⁶ We also have a deep embedding for all logics. We use the shallow embedding as it is much briefer and clearer: the interaction between the assertion logic and P3TL, although non-trivial, is not the main focus of this paper.

Note that our definition of the semantics of `Secure` ensures that programs cannot get stuck because of policy violation. The only reason that there might be no s' such that $s \xrightarrow{-e} s'$ in this theorem is non-termination of while loops.

5 Synthesis

In this section we discuss the synthesis algorithm; that is, the algorithm that takes a P3TL formula and emits a reference monitor which enforces the policy, and a proof that the generated monitor does indeed enforce the policy.

We will use the example policy

$$(\ominus \langle \lambda s. x \ s = 1 \rangle) [\Rightarrow] (\langle \lambda s. x \ s < 5 \rangle \ \mathcal{S} \ \langle \lambda s. y \ s = 1 \rangle) \quad (*)$$

as a running example throughout this section. Although this example has no correspondence to a real-world policy (even a small such policy would be too big for the limited space available), it contains enough complexity to be of interest.

5.1 Checking P3TL satisfaction

In the synthesis of P3TL formulae, the following definition is required

Definition 1. *The past formulae of a formula are sub-formulae*

1. of the form $\phi \ \mathcal{S} \ \psi$; or
2. that occur after \ominus , e.g. ψ in $\ominus \psi$.

The number of sub-formulae is linear in the size of the formula.

Example 5. The past formulae for our example policy (*) are $\langle \lambda s. x \ s = 1 \rangle$ and $\langle \lambda s. x \ s < 5 \rangle \ \mathcal{S} \ \langle \lambda s. y \ s = 1 \rangle$. Terms in our examples which relate to these formulae will have suffixes 0 and 1, respectively. For example, `state_0` is $\Sigma \langle \lambda s. x \ s = 1 \rangle$.

Note that a since formula may be unfolded according to the equality:

$$\phi \ \mathcal{S} \ \psi = \psi \ [\vee] (\ominus (\phi \ \mathcal{S} \ \psi) \ [\wedge] \ \phi)$$

Applying this rule repeatedly yields a formula in which all since formulae occur only after a \ominus operator. Thus, truth of the rewritten formula depends only on the truth of propositions in the current world and the truth of sub-formulae in the immediately previous world. Furthermore, note that these sub-formulae are all past formulae.

To check if a sequence σ satisfies a policy, we start by setting all formulae of the form $\ominus \ \varphi$ to false (as per the semantics of \ominus). Starting from the initial state, check and record the truth of all past formulae. At the last state in σ , check the truth of the rewritten policy, using the recorded past formulae from the previous state.

Thus, to monitor a policy, we need only keep track of its past formulae. This implies that we do not require an explicit representation of worlds in our monitor, and that checking a P3TL formula can be done efficiently.

```

IF ( $\lambda s. x\ s = 1$ ) THEN
  Do ( $\lambda s. s(\text{tmp}_0 := 1)$ )
ELSE
  Do ( $\lambda s. s(\text{tmp}_0 := 0)$ )
FI
IF ( $\lambda s. y\ s = 1$ ) THEN
  Do ( $\lambda s. s(\text{tmp}_1 := 1)$ )
ELSE
  IF ( $\lambda s. \text{state}_1\ s = 1$ ) THEN
    IF ( $\lambda s. x\ s < 5$ ) THEN
      Do ( $\lambda s. s(\text{tmp}_1 := 1)$ )
    ELSE
      Do ( $\lambda s. s(\text{tmp}_1 := 0)$ )
    FI
  ELSE
    Do ( $\lambda s. s(\text{tmp}_1 := 0)$ )
  FI
FI

```

(CONT.)

```

IF ( $\lambda s. \text{state}_0\ s = 1$ ) THEN
  IF ( $\lambda s. \text{tmp}_1\ s = 1$ ) THEN
    Secure policy;
    Do ( $\lambda s. s(\text{state}_0 := \text{tmp}_0\ s)$ );
    Do ( $\lambda s. s(\text{state}_1 := \text{tmp}_1\ s)$ )
  ELSE
    Skip
  FI
ELSE
  Secure policy;
  Do ( $\lambda s. s(\text{state}_0 := \text{tmp}_0\ s)$ );
  Do ( $\lambda s. s(\text{state}_1 := \text{tmp}_1\ s)$ )
FI

```

Fig. 5. Generated code for policy (*). The left hand column contains the state maintenance code, while the right side contains the policy checking code.

5.2 Monitor synthesis

The algorithm for constructing a monitor for a P3TL formula is then as follows:

1. For each past formulae ψ , allocate a state bit, Σ_ψ , which records the truth of the formula in the previous world, i.e. $\Sigma_\psi \leftrightarrow \ominus\psi$;
2. Construct a program fragment which checks the truth of each past formula ψ in the current world, i.e., with respect to the current program state. When constructing the fragment, if a sub-formula of ψ of the form $\ominus\phi$ is seen, emit code which checks the state bit Σ_ϕ ;
3. Construct a monitor fragment for the main formula. In the case that the formula holds, execute the secure statement and update the monitor state, otherwise handle the security violation; and
4. Sequentially compose the fragments to generate the final monitor.

Fig. 6 presents the algorithm for constructing a monitor fragment. The notation $\mathcal{S}[\psi]_{(tc,fc)}$ denotes the algorithm applied to formula ψ , with arguments **tc** (true case) and **fc** (false case). The leaves of the fragment, **tc** and **fc**, are assignments in the case of a past formula, and a secure statement along with state update in the case of the policy.

Example 6. The monitor generated for our example policy is shown in Fig. 5. Our implementation includes an optimisation: rather than re-check a formula, monitor fragments may refer to previously established past formula by checking the corresponding variable.

$$\begin{aligned}
\mathcal{S}[\llbracket a \rrbracket]_{(tc, fc)} &= \text{IF } a \text{ THEN } (tc \text{ ATOMI}) \text{ ELSE } (fc \text{ NATOMI}) \\
\mathcal{S}[\llbracket \psi \Rightarrow \phi \rrbracket]_{(tc, fc)} &= \text{let} \\
&\quad tc' = \lambda r. \mathcal{S}[\llbracket \phi \rrbracket]_{(tc \cdot \text{IMPLI}_1, fc \cdot (\text{NIMPLI } r))} \\
&\quad \text{in} \\
&\quad \mathcal{S}[\llbracket \psi \rrbracket]_{(tc', tc \cdot \text{IMPLI}_2)} \\
\mathcal{S}[\llbracket \neg \psi \rrbracket]_{(tc, fc)} &= \mathcal{S}[\llbracket \psi \rrbracket]_{(fc \cdot \text{NNEGI}, tc \cdot \text{NEGI})} \\
\mathcal{S}[\llbracket \ominus \psi \rrbracket]_{(tc, fc)} &= \text{IF } (\Sigma_\psi) \text{ THEN } (tc \text{ INV-}\pi_\psi) \text{ ELSE } (fc \text{ NINV-}\pi_\psi) \\
\mathcal{S}[\llbracket \psi \mathcal{S} \phi \rrbracket]_{(tc, fc)} &= \mathcal{S}[\llbracket \phi \vee (\ominus(\psi \mathcal{S} \phi) \wedge \psi) \rrbracket]_{(tc \cdot \text{SINCEI}, fc \cdot \text{NSINCEI})}
\end{aligned}$$

Fig. 6. The algorithm for constructing a monitor fragment. It generates both a monitor fragment and proof annotations. The parameters tc and fc are functions which take proof annotations and return the statements to be used in each branch of the conditional; the term $tc \cdot \text{SINCEI}$ composes rule SINCEI with tc , thus adding a new rule to the proof tree at tc .

$s \models \varphi$

$$\begin{array}{c}
\frac{\varphi \text{ (fst } s)}{s \models \llbracket \varphi \rrbracket} \text{ATOMI} \quad \frac{\neg \varphi \text{ (fst } s)}{\neg s \models \llbracket \varphi \rrbracket} \text{NATOMI} \\
\frac{s \models \varphi}{s \models \psi \Rightarrow \varphi} \text{IMPLI}_1 \quad \frac{\neg s \models \psi}{s \models \psi \Rightarrow \varphi} \text{IMPLI}_2 \quad \frac{s \models \varphi \quad \neg s \models \psi}{\neg s \models \varphi \Rightarrow \psi} \text{NIMPLI} \\
\frac{\neg s \models \varphi}{s \models \neg \varphi} \text{NEGI} \quad \frac{s \models \varphi}{\neg s \models \neg \varphi} \text{NNEGI} \\
\frac{s \models \psi \vee (\ominus(\varphi \mathcal{S} \psi) \wedge \varphi)}{s \models \varphi \mathcal{S} \psi} \text{SINCEI} \\
\frac{\neg s \models \psi \vee (\ominus(\varphi \mathcal{S} \psi) \wedge \varphi)}{\neg s \models \varphi \mathcal{S} \psi} \text{NSINCEI}
\end{array}$$

Fig. 7. Derived introduction rules for P3TL.

The monitor fragment for the past formula $\llbracket \lambda s. x \ s = 1 \rrbracket$ (the first 5 lines of Fig. 5) was generated by

$$\mathcal{S}[\llbracket \lambda s. x \ s = 1 \rrbracket]_{(\text{Do } (\lambda s. s(\text{tmp}_0 := 1)), \text{Do } (\lambda s. s(\text{tmp}_0 := 0)))}$$

Not shown are the proof handling aspects; in our prototype, the Do statements are functions taking a proof of $\llbracket \lambda s. x \ s = 1 \rrbracket$ and $\neg \llbracket \lambda s. x \ s = 1 \rrbracket$ respectively. These proofs are then stored for later extraction.

If we, for the moment, ignore proof generation, the synthesis algorithm is straightforward: when an atom is seen, generate code which checks the atom and uses tc in the true case, fc in the false case; in the case of a negation, $\neg \psi$, generate code for ψ but switch the leaves — if ψ holds, then use fc , otherwise tc . Note that negation produces no extra code; it merely swaps leaves. This means that special cases for conjunction and disjunction are not required: unfolding

the definitions results in no additional code; a previous formula, $\ominus\psi$, results in a check of the state component for that formula, Σ_ψ .

5.3 Proof synthesis

The main monitor theorem is

$$\Vdash \{\text{INV}_\psi\} \text{ monitor } \{\text{INV}_\psi\}$$

which states that the monitor preserves the invariant, discussed below. The proof that the monitor enforces the security policy is implicit in the use of any **Secure** statements: wherever an operation occurs in the monitor, a proof obligation is required which states that the policy holds (c.f. rule **SECURE**). Thus, as our proof system is sound (Theorem 1), a proof of the main theorem implies that execution of a **Secure** statement occurs only when the security policy is true.

The proof of the main theorem is generated from the monitor fragments using a verification condition generator style algorithm. The proof requires a number of lemmas for both re-establishing the monitor invariant and for showing the policy holds for **Secure** statements. The generation of these proof obligations is discussed below.

The monitor invariant The monitor maintains state between invocations, in particular Σ_ϕ for each past formula ϕ , and so a monitor invariant is required. This invariant relates the value of each state variable to the truth of the corresponding formula in the previous state.

$$\text{INV}_\psi \equiv \bigwedge_{\phi \in \text{past}(\psi)} \Sigma_\phi \leftrightarrow \ominus\phi$$

where ψ is the security policy and $\text{past}(\psi)$ are the past formulae of ψ .

The past formula monitor fragments are then responsible for reestablishing the monitor invariant. This is done in two steps: firstly, the monitor checks the past formula, ϕ , and sets a temporary variable, Δ_ϕ , accordingly; secondly, the monitor updates the real monitor state *after* execution of the **Secure** statement.

This two-step process is required for a number of reasons: primarily, later monitor fragments may require the state variable to establish the truth of other formulae, past or policy; and, secondly, if the security policy doesn't hold, then the monitor may elect to silently ignore the request. In this case, the initial value of the state variables is still correct — no secure operation is performed, and thus the invariant is still true. This is the behaviour of monitors generated by our prototype. We then define, for each past formula, a pre-invariant $\Xi_\phi \equiv \Delta_\phi \leftrightarrow \phi$. This correspondence is used to generate the invariant — after execution of **Secure** statement and state update, we can use this to derive $\Sigma_\phi \leftrightarrow \ominus\phi$

Example 7. Our prototype generates a number of auxiliary lemmas for manipulating the invariant. These include

- projection rules (the $\text{INV-}\pi_\phi$ rules mentioned in Fig. 6)

$$\frac{\text{invariant } (s, \sigma) \quad \text{state_1 } s = 1}{(s, \sigma) \models \ominus (\llbracket \lambda s. x \ s < 5 \rrbracket \ S \ \llbracket \lambda s. y \ s = 1 \rrbracket)} \text{INV-}\pi_1$$

- rules for invariance under assignment (X is an arbitrary function)

$$\text{invariant } (s(\text{tmp_1} := X \ s), \sigma) = \text{invariant } (s, \sigma)$$

Proof obligations The proof of the main theorem requires, for each assignment in each past formula monitor fragment, a proof obligation of the form

$$\text{INV}_\psi \wedge \Xi_{\phi_1} \wedge \dots \wedge \Xi_{\phi_m} \wedge a_1 \wedge \dots \wedge a_n \longrightarrow \text{INV}'_\psi \wedge \Xi'_{\phi_1} \wedge \dots \wedge \Xi'_{\phi_m} \wedge \Xi'_\phi$$

A formula is primed to denote it's truth after the assignment. The terms $\Xi_{\phi_1} \wedge \dots \wedge \Xi_{\phi_m}$ are the past formula equivalences established by previous monitor fragments, a_1, \dots, a_n are the atoms (or their negation, in the case of a false branch) that were checked by the conditionals in the current monitor fragment, and ϕ is the formula from which the monitor fragment was produced, with Δ_ϕ replaced by **True** or **False** depending on which branch the assignment occurs.

These proofs state that the monitor fragment correctly establishes Ξ_ϕ and does not invalidate previously established equivalences or the invariant.

The obligations for **Secure** statements are similar but for the last term in the conjunction: in this case, it is the policy. In addition, the **SECURE** rule allows us to generate P3TL terms of the form $\ominus\phi$, assuming we have ϕ — this is how the conversion from pre-invariant to invariant occurs.

Example 8. The following obligation is generated for the fourth assignment in Fig. 5.

$$\begin{aligned} \forall s \ \sigma. \text{invariant } (s, \sigma) \wedge \\ \text{pre_sd0 } (s, \sigma) \wedge y \ s \neq 1 \wedge \text{state_1 } s = 1 \wedge x \ s < 5 \longrightarrow \\ \text{invariant } (s(\text{tmp_1} := 1), \sigma) \wedge \\ \text{pre_sd0 } (s(\text{tmp_1} := 1), \sigma) \wedge \text{pre_sd1 } (s(\text{tmp_1} := 1), \sigma) \end{aligned}$$

Proof construction Much of complexity of the algorithm in Fig. 6 arises because the leaves of the tree are annotated with the proof, and the algorithm builds the tree from the leaves up. The proofs require information that is not initially available, i.e., the proofs for sub-formulae, so the arguments to the synthesis function, **tc** and **fc**, are *functions* from proofs to program fragments. In particular, if the fragment is checking the truth of ϕ , **tc** will be a function from a proof of ϕ , and **fc** a function from a proof of $\neg\phi$.

Fig. 7 shows the P3TL proof rules used by the algorithm, except for the invariant projection lemmas which are described above. All of these rules are derived from the semantics of P3TL (using Isabelle) as lemmas, but are intended to be derivable in a syntactic proof system (such as that in Lichtenstein and Pnueli [11]). Each rule occurs in the positive and negated form, as some proof

we require temporal terms only when dealing with the high-level safety policy; this should make extending existing programs simpler. Nevertheless, we envisage no major issues in synthesising monitors for their framework.

Synthesis from temporal logics, traditionally used in the model checking community (e.g. SPIN [10]) has gained recent popularity for constructing program reference monitors [6,7,2].

In particular, our approach is similar to that of the `PATHEXPLORER` project [9]. They construct monitors for safety properties using an algorithm that is very close to that we presented. These monitors can then be automatically inserted into Java programs for run-time testing. The major difference is that our algorithm also generates proofs.

Peled and Zuck [19] generate a proof from model-checking results. This proof shows properties of the target system; in theory, we may be able to generate the monitor and then apply their technique to generate the proof. It is, however, unclear whether that approach would be feasible in practice.

7 Conclusions and Future Work

In this paper we introduced a temporal logic for formulating security policies — propositional pure past temporal logic — and showed how to automatically generate efficient reference monitor implementations that check the required policy, along with a machine-checkable proof of their safety.

We have implemented a prototype targeting Isabelle/HOL; the majority of formal matter, and all of the theorems, in this paper were generated using Isabelle’s presentation mechanism [18] from the Isabelle proofs. This means what we show is what we proved.

The main contribution of this paper is to show how reference monitor synthesis, proof generation, and the policy logic are defined and interact. In future, we are interested in refining the system towards one in which the reference monitors are implemented in a low-level language and inserted into consumer code by binary rewriting, as demonstrated in [21]. Also desirable is a richer policy logic, such as a first order variant of P3TL; finally, a higher-level language that can be compiled into P3TL would ease the job of writing security policies.

Acknowledgements

We wish to thank Kai Engelhardt and Harvey Tuch, along with the anonymous reviewers, for their valuable comments on earlier versions of this paper.

References

1. D. Aspinall, S. Gilmore, M. Hofmann, D. Sannella, and I. Stark. Mobile resource guarantees for smart devices. In *Proc. of CASSIS’04*, volume 3362 of *LNCS*, pages 1–26. Springer, Mar. 2005.

2. H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Program monitoring with LTL in EAGLE. In *Proc. of PADTAD'04*, Apr. 2004.
3. G. Barthe, T. Rezk, and A. Saabas. Proof obligations preserving compilation. In *Proc. of FAST'05*, volume 3866 of *LNCS*, pages 112–126. Springer, July 2006.
4. A. Bernard and P. Lee. Temporal logic for proof-carrying code. In *Proc. of CADE'02*, pages 31–46, London, UK, 2002. Springer-Verlag.
5. D. F. C. Brewer and M. J. Nash. The Chinese Wall security policy. In *IEEE Symposium on Security and Privacy*, pages 206–214, 1989.
6. F. Chen, M. d'Amorim, and G. Rosu. A formal monitoring-based framework for software development and analysis. In *Proc. of ICFEM'04*, volume 3308 of *LNCS*, pages 357–372. Springer, Nov. 2004.
7. M. d'Amorim and G. Rosu. Efficient monitoring of omega-languages. volume 3576 of *LNCS*, pages 364–378, July 2005.
8. M. Geilen. On the construction of monitors for temporal logic properties. volume 55, 2001.
9. K. Havelund and G. Rosu. Synthesizing monitors for safety properties. In *Proc. of TACAS'02*, volume 2280 of *LNCS*, pages 342–356. Springer, Apr. 2002.
10. G. J. Holzmann. The model checker spin. *IEEE Trans. Software Eng.*, 23(5):279–295, 1997.
11. O. Lichtenstein and A. Pnueli. Propositional temporal logics: Decidability and completeness. *Logic Journal of the IGPL*, 8(1):55–85, 2000.
12. Z. Manna and A. Pnueli. The anchored version of the temporal framework. In J. W. de Bakker, W. P. de Roever, and G. Rozenberg, editors, *Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency*, volume 354 of *LNCS*, pages 201–284. Springer, 1989.
13. Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer-Verlag, New York, 1995.
14. G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly language. *TOPLAS*, 21(3):527–568, 1999.
15. G. C. Necula. Proof-carrying code. In *Proc. of POPL'97*, pages 106–119, Paris, Jan. 1997.
16. G. C. Necula and P. Lee. The design and implementation of a certifying compiler. In *Proc. of PLDI'98*, volume 33,5, pages 333–344, New York, June 17–19 1998. ACM Press.
17. T. Nipkow. Hoare logics in Isabelle/HOL. In H. Schwichtenberg and R. Steinbrüggen, editors, *Proof and System-Reliability*, pages 341–367. Kluwer, 2002.
18. T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
19. D. Peled and L. Zuck. From model checking to a temporal proof. In *Proc. of SPIN'01*, pages 1–14, New York, NY, USA, 2001. Springer.
20. F. B. Schneider. Enforceable security policies. *Information and System Security*, 3(1):30–50, 2000.
21. S. Winwood and M. M. T. Chakravarty. Secure untrusted binaries - provably!. In *Proc. of FAST'05*, volume 3866 of *LNCS*, pages 171–186. Springer, 2006.