

The Spirit of Ghost Code

Jean-Christophe Filliâtre^{1,2}, Léon Gondelman^{1*}, and Andrei Paskevich^{1,2}

¹ Lab. de Recherche en Informatique, Univ. Paris-Sud, CNRS, Orsay, F-91405

² INRIA Saclay – Île-de-France, Orsay, F-91893

Abstract. In the context of deductive program verification, ghost code is part of the program that is added for the purpose of specification. Ghost code must not interfere with regular code, in the sense that it can be erased without observable difference in the program outcome. In particular, ghost data cannot participate in regular computations and ghost code cannot mutate regular data or diverge. The idea exists in the folklore since the early notion of auxiliary variables and is implemented in many state-of-the-art program verification tools. However, a rigorous definition and treatment of ghost code is surprisingly subtle and few formalizations exist.

In this article, we describe a simple ML-style programming language with mutable state and ghost code. Non-interference is ensured by a type system with effects, which allows, notably, the same data types and functions to be used in both regular and ghost code. We define the procedure of ghost code erasure and we prove its safety using bisimulation. A similar type system, with numerous extensions which we briefly discuss, is implemented in the program verification environment Why3.

1 Introduction

A common technique in deductive program verification consists in introducing data and computations, traditionally named *ghost code*, that only serve to facilitate specification. Ghost code can be safely erased from a program without affecting its final result. Consequently, a ghost expression cannot be used in a *regular* (non-ghost) computation, it cannot modify a regular mutable value, and it cannot raise exceptions that would escape into regular code. However, a ghost expression can use regular values and its result can be used in program annotations: preconditions, postconditions, loop invariants, assertions, etc. A classical use case for ghost code is to equip a data structure with ghost fields containing auxiliary data for specification purposes. Another example is ghost step counters to prove the time complexity of an algorithm.

When it comes to computing verification conditions, for instance using a weakest precondition calculus, there is no need to make a distinction between ghost and regular code. At this moment, ghost code is just a computation that supplies auxiliary values to use in specification and to simplify proofs. This

* This work is partly supported by the Bware (ANR-12-INSE-0010, <http://bware.lri.fr/>) project of the French national research organization (ANR).

computation, however, is not necessary for the program itself and thus should be removed when we compile the annotated source code. Therefore we need a way to ensure, by static analysis, that ghost code does not interfere with the rest of the program.

Despite that the concept of ghost code exists since the early days of deductive program verification, and is supported in most state-of-the-art tools [1–4], it is surprisingly subtle. In particular, a sound non-interference analysis must ensure that every ghost sub-expression terminates. Otherwise, one could supply such a sub-expression with an arbitrary postcondition and thus be able to prove anything about the program under consideration. Another non-obvious observation is that structural equality cannot be applied naively on data with ghost components. Indeed, two values could differ only in their ghost parts and consequently the comparison would yield a different result after the ghost code erasure.

There is a number of design choices that show up when conceiving a language with ghost code. First, how explicit should we be in our annotations? For example, should every ghost variable be annotated as such, or can we infer its status by looking at the values assigned to it? Second, how much can be shared between ghost and regular code? For instance, can a ghost value be passed to a function that does not specifically expect a ghost argument? Similarly, can we store a ghost value in a data structure that is not specifically designed to hold ghost data, *e.g.* an array or a tuple? Generally speaking, we should decide where ghost code can appear and what can appear in ghost code.

In this article, we show that, using a tailored type system with effects, we can design a language with ghost code that is both expressive and concise. As a proof of concept, we describe a simple ML-style programming language with mutable state, recursive functions, and ghost code. Notably, our type system allows the same data types and functions to be used in both regular and ghost code. We give a formal proof of the soundness of ghost code erasure, using a bisimulation argument. A type system based on the same concepts is implemented in the verification tool Why3 [4]. The language presented in this paper is deliberately simplified. The more exciting features, listed in Section 4 and implemented in Why3, only contribute to more complex effect tracking in the type system, which is mostly orthogonal to the problem of ghost code non-interference.

Outline. This paper is organized as follows. Section 2 introduces an ML-like language with ghost code. Section 3 defines the operation of ghost code erasure and proves its soundness. Section 4 describes the actual implementation in Why3. We conclude with related work in Section 5 and perspectives in Section 6. An extended version of this paper containing proofs is available at <http://hal.archives-ouvertes.fr/hal-00873187/>.

2 GhostML

We introduce GhostML, a mini ML-like language with ghost code. It features global references (that is, mutable variables), recursive functions, and integer and Boolean primitive types.

2.1 Syntax

The syntax of GhostML is given in Fig. 1. Terms are either values or compound expressions like application, conditional, reference access and modification. We assume a fixed finite set of global references. All the language constructions are standard ML, except for the keyword **ghost** which turns a term t into ghost code.

$t ::=$		TERMS	$v ::=$		VALUES
$ $	v	<i>value</i>	$ $	c	<i>constant</i>
	$t v$	<i>application</i>		x^β	<i>variable</i>
	$\text{let } x^\beta = t \text{ in } t$	<i>local binding</i>		$\lambda x^\beta : \tau. t$	<i>anonymous function</i>
	$\text{if } v \text{ then } t \text{ else } t$	<i>conditional</i>		$\text{rec } x^\beta : \tau^\beta \Rightarrow \tau. \lambda x^\beta : \tau. t$	<i>recursive function</i>
	$r^\beta := v$	<i>assignment</i>			
	$!r^\beta$	<i>dereference</i>			
$\tau ::=$	$ $		$c ::=$	$ $	CONSTANTS
	κ	<i>primitive type</i>		$()$	<i>unit</i>
	$ $			$\dots, -1, 0, 1, \dots$	<i>integers</i>
	$\tau^\beta \Rightarrow \tau$	<i>functional type</i>		$\text{true}, \text{false}$	<i>Boolean</i>
$\kappa ::=$	$ $	PRIMITIVE TYPES	$\beta \in \{\perp, \top\}$	$ $	GHOST STATUS
	$\text{int} \mid \text{bool} \mid \text{unit}$	<i>primitive types</i>		$\epsilon \in \{\perp, \top\}$	EFFECT

Fig. 1. Syntax.

Every variable is tagged with a ghost status β , which is $\top$ for ghost variables and $\perp$ for regular ones (here and below, “regular” stands for “non-ghost”). Similarly, references and formal function parameters carry their ghost status. Consider the following example:

$$\text{let } \text{upd}^\top = \lambda x^\perp : \text{int}. g^\top := x^\perp \text{ in } \text{upd}^\top !r^\perp$$

Here, function $\text{upd}^\top$ takes one regular parameter $x^\perp$ and assigns it to a ghost reference $g^\top$. Then $\text{upd}^\top$ is applied to the contents of a regular reference $r^\perp$.

Note that compound terms obey a variant of *A-normal* form [5]. That is, in application, conditional, and reference assignment, one of the sub-expressions must be a value. This does not reduce expressiveness, since a term such as $(t_1 (t_2 v))$ can be rewritten as $\text{let } x^\beta = t_2 v \text{ in } t_1 x^\beta$, where β depends on the ghost status of the first formal parameter of t_1 .

Types are either primitive data-types (int , bool , unit) or function types. A function type is an arrow $\tau_2^\beta \Rightarrow \tau_1$ where β stands for the function argument’s ghost status, and ϵ is the *latent* effect of the function. An effect ϵ is a Boolean value that indicates presence of regular side effects such as modification of a regular reference or possible non-termination.

MiniML Syntax. The syntax of traditional MiniML can be obtained by omitting all ghost indicators β (on references, variables, parameters, and types) and excluding the `ghost` construct. Equivalently, we could define MiniML as the subset of GhostML where all ghost indicators β are $\perp$ and where terms of the form `ghost` t do not appear.

2.2 Semantics

Fig. 2 gives a small-step operational semantics to GhostML which corresponds to a deterministic call-by-value reduction strategy. Each reduction step defines a relation between states. A state is a pair $t \mid \mu$ of a term t and a store μ . A store μ maps global references of t to constants. The regular part of a store μ , written $\mu_{\perp}$, is the restriction of μ to regular references. Rules indicate the store μ only when relevant.

A reduction step can take place directly at the top of a term t . Such a step is called a head reduction and is denoted $t \mid \mu \xrightarrow{\alpha} t' \mid \mu'$. Rule (E-GHOST) expresses that, from the point of view of operational semantics, there is no difference between regular and ghost code. Other head reduction rules are standard. For instance, rules (E-OP- λ) and (E-OP- δ) evaluate the application of a constant c_0 to constants $c_1 \dots c_m$. Such an application is either partial ($1 \leq m < \text{arity}(c_0)$), and then turned into a function $\lambda x^{\perp} : \kappa. c_0 \ c_1 \ \dots \ c_m \ x^{\perp}$, or total ($m = \text{arity}(c_0)$), and then some oracle function δ gives the result $\delta(c_0, c_1, \dots, c_m)$. For instance, $\delta(\text{not}, \text{true}) = \text{false}$, $\delta(+, 47, -5) = 42$, etc.

A reduction step can also be contextual, *i.e.* it takes place in some sub-expression. Since our language is in A-normal form, there are only two contextual rules, (E-CONTEXT-APP) and (E-CONTEXT-LET).

As usual, $\rightarrow^*$ denotes the reflexive, transitive closure of $\rightarrow$. We say that a closed term t evaluates to v in a store μ if there is a μ' such that $t \mid \mu \rightarrow^* v \mid \mu'$. Note that, since t is closed, v is not a variable. Finally, the divergence of a term t in a store μ is defined co-inductively as follows:

$$\frac{t \mid \mu \rightarrow^1 t' \mid \mu' \quad t' \mid \mu' \rightarrow \infty}{t \mid \mu \rightarrow \infty} \text{(E-Div)}$$

MiniML Semantics. Since ghost statuses do not play any role in the semantics of GhostML, dropping them (or, equivalently, marking all β as $\perp$) and removing the rule (E-GHOST) results in a standard call-by-value small-step operational semantics for MiniML. For the sake of clarity, we use a subscript m when writing MiniML reduction steps: $t \mid \mu \rightarrow_m t' \mid \mu'$.

2.3 Type System

The purpose of the type system is to ensure that “well-typed terms do not go wrong”. In our case, “do not go wrong” means not only that well-typed terms verify the classical type soundness property, but also that ghost code *does not*

$\text{ghost } t \xrightarrow{\alpha} t$	(E-GHOST)
$\frac{1 \leq m < \text{arity}(c_0)}{c_0 \ c_1 \ \dots \ c_m \xrightarrow{\alpha} \lambda x^{\perp} : \kappa. c_0 \ c_1 \ \dots \ c_m \ x^{\perp}}$	(E-OP- λ)
$\frac{m = \text{arity}(c_0) \quad \delta(c_0, c_1, \dots, c_m) \text{ is defined}}{c_0 \ c_1 \ \dots \ c_m \xrightarrow{\alpha} \delta(c_0, c_1, \dots, c_m)}$	(E-OP- δ)
$(\lambda x^{\beta} : \tau. t) \ v \xrightarrow{\alpha} t[x^{\beta} \leftarrow v]$	(E-APP- λ)
$(\text{rec } f^{\beta} : \tau^{\beta} \Rightarrow \tau. \lambda x^{\beta} : \tau. t) \ v \xrightarrow{\alpha} t[x^{\beta} \leftarrow v, f^{\beta} \leftarrow \text{rec } f^{\beta} : \tau^{\beta} \Rightarrow \tau. \lambda x^{\beta} : \tau. t]$	(E-APP-REC)
$\text{let } x^{\beta} = v_1 \text{ in } t_2 \xrightarrow{\alpha} t_2[x^{\beta} \leftarrow v_1]$	(E-LET)
$\text{if true then } t_1 \text{ else } t_2 \xrightarrow{\alpha} t_1$	(E-IF-TRUE)
$\text{if false then } t_1 \text{ else } t_2 \xrightarrow{\alpha} t_2$	(E-IF-FALSE)
$!r^{\beta} \mid \mu \xrightarrow{\alpha} \mu(r^{\beta}) \mid \mu$	(E-DEREF)
$r^{\beta} := c \mid \mu \xrightarrow{\alpha} () \mid \mu[r^{\beta} \mapsto c]$	(E-ASSIGN)
$\frac{t \mid \mu \xrightarrow{\alpha} t' \mid \mu'}{t \mid \mu \rightarrow t' \mid \mu'}$	(E-HEAD)
$\frac{t_1 \mid \mu \rightarrow t'_1 \mid \mu'}{(t_1 \ v) \mid \mu \rightarrow (t'_1 \ v) \mid \mu'}$	(E-CONTEXT-APP)
$\frac{t_2 \mid \mu \rightarrow t'_2 \mid \mu'}{\text{let } x^{\beta} = t_2 \text{ in } t_1 \mid \mu \rightarrow \text{let } x^{\beta} = t'_2 \text{ in } t_1 \mid \mu'}$	(E-CONTEXT-LET)

Fig. 2. Semantics.

interfere with regular code. More precisely, non-interference means that ghost code never modifies regular references and that it always terminates. For that purpose, we introduce a type system with effects, where the typing judgment is

$$\Sigma, \Gamma \vdash t : \tau, \beta, \epsilon.$$

Here, τ is the type of term t . Boolean indicators β and ϵ indicate, respectively, the ghost status of t and its regular side effects. Γ is a typing environment that binds variables to types. Σ is a store typing that binds each reference r^{β} to the primitive type of the stored value. We restrict types of stored values to primitive types to avoid a possible non-termination via Landin's knot (that is,

$\frac{\text{Typeof}(c) = \tau}{\Sigma, \Gamma \vdash c : \tau, \perp, \perp}$	(T-CONST)
$\frac{(x^\beta : \tau) \in \Gamma}{\Sigma, \Gamma \vdash x^\beta : \tau, \beta, \perp}$	(T-VAR)
$\frac{\Sigma, \Gamma, x^\beta : \tau \vdash t : \tau_0, \beta_0, \epsilon}{\Sigma, \Gamma \vdash (\lambda x^\beta : \tau. t) : \tau^\beta \Rightarrow \tau_0, \beta_0, \perp}$	(T- λ)
$\frac{\Sigma, \Gamma, f^\perp : \tau_2^\beta \Rightarrow \tau_1 \vdash (\lambda x^\beta : \tau_2. t) : \tau_2^\beta \Rightarrow \tau_1, \perp, \perp}{\Sigma, \Gamma \vdash (\text{rec } f^\perp : \tau_2^\beta \Rightarrow \tau_1. \lambda x^\beta : \tau_2. t) : \tau_2^\beta \Rightarrow \tau_1, \perp, \perp}$	(T-REC)
$\frac{\Sigma, \Gamma \vdash v : \text{bool}, \beta_0, \perp \quad \Sigma, \Gamma \vdash t_1 : \tau, \beta_1, \epsilon_1 \quad \Sigma, \Gamma \vdash t_2 : \tau, \beta_2, \epsilon_2}{\Sigma, \Gamma \vdash (\text{if } v \text{ then } t_1 \text{ else } t_2) : \tau, \beta_0 \vee \beta_1 \vee \beta_2, \epsilon_1 \vee \epsilon_2}$	(T-IF)
$\frac{\Sigma, \Gamma, x^\perp : \tau_2 \vdash t_1 : \tau_1, \beta_1, \epsilon_1 \quad \Sigma, \Gamma \vdash t_2 : \tau_2, \beta_2, \epsilon_2}{\Sigma, \Gamma \vdash (\text{let } x^\perp = t_2 \text{ in } t_1) : \tau_1, \beta_1 \vee \beta_2, \epsilon_1 \vee \epsilon_2}$	(T-LET-REGULAR)
$\frac{\Sigma, \Gamma, x^\top : \tau_2 \vdash t_1 : \tau_1, \beta_1, \epsilon_1 \quad \Sigma, \Gamma \vdash t_2 : \tau_2, \beta_2, \perp}{\Sigma, \Gamma \vdash (\text{let } x^\top = t_2 \text{ in } t_1) : \tau_1, \beta_1, \epsilon_1}$	(T-LET-GHOST)
$\frac{\Sigma, \Gamma \vdash t : \tau_2^\perp \Rightarrow \tau_1, \beta_1, \epsilon_2 \quad \Sigma, \Gamma \vdash v : \tau_2, \beta_2, \perp}{\Sigma, \Gamma \vdash (t \ v) : \tau_1, \beta_1 \vee \beta_2, \epsilon_1 \vee \epsilon_2}$	(T-APP-REGULAR)
$\frac{\Sigma, \Gamma \vdash t : \tau_2^\top \Rightarrow \tau_1, \beta_1, \epsilon_2 \quad \Sigma, \Gamma \vdash v : \tau_2, \beta_2, \perp}{\Sigma, \Gamma \vdash (t \ v) : \tau_1, \beta_1, \epsilon_1 \vee \epsilon_2}$	(T-APP-GHOST)
$\frac{(r^\beta : \kappa) \in \Sigma}{\Sigma, \Gamma \vdash !r^\beta : \kappa, \beta, \perp}$	(T-DEREF)
$\frac{\Sigma, \Gamma \vdash v : \kappa, \beta', \perp \quad (r^\beta : \kappa) \in \Sigma \quad \beta \geq \beta'}{\Sigma, \Gamma \vdash (r^\beta := v) : \text{unit}, \beta, \neg\beta}$	(T-ASSIGN)
$\frac{\Sigma, \Gamma \vdash t : \tau, \beta, \perp}{\Sigma, \Gamma \vdash (\text{ghost } t) : \tau, \top, \perp}$	(T-GHOST)

Fig. 3. Typing rules.

recursion encoded using a mutable variable containing a function), which would be undetected in our type system.

Typing rules are given in Fig. 3. To account for non-interference, each rule whose conclusion is a judgement $\Sigma, \Gamma \vdash t : \tau, \beta, \epsilon$ is added the implicit extra side condition

$$(\beta = \top) \Rightarrow (\epsilon \vee \epsilon^+(\tau) = \perp) \quad (1)$$

where $\epsilon^+(\tau)$ is defined recursively on τ as follows:

$$\begin{aligned} \epsilon^+(\kappa) &\triangleq \perp \\ \epsilon^+(\tau_2 \xRightarrow{\beta} \tau_1) &\triangleq \epsilon \vee \epsilon^+(\tau_1) \end{aligned}$$

In other words, whenever t is ghost code, it must terminate and must not modify any regular reference. In particular, a ghost function whose body is possibly non-terminating or possibly modifies a regular reference is rejected by the type system.

Let us explain some rules in detail. The rule (T-CONST) states that any constant c is regular code, (i.e. $\beta = \perp$) yet is pure and terminating (i.e. $\epsilon = \perp$). Moreover, we assume that if c is some constant operation, then its formal parameters are all regular. The type of each constant is given by some oracle function $\text{Typeof}(c)$. For instance, $\text{Typeof}(+) = \text{int}^{\perp} \multimap \text{int}^{\perp} \multimap \text{int}$.

Recursive functions are typed as follows. For simplicity, we assume that whenever a recursive function is used, we may have non-termination. Therefore, we enforce the latent effect ϵ of any recursive function to be $\top$. Consequently, no recursive function can be used or even occur in ghost code. In practice, however, we do not have to assign a latent non-termination effect to recursive functions whose termination can be established by static analysis (e.g. by a formal proof).

The rule (T-IF) shows how ghost code is propagated through conditional expressions: if at least one of the branches or the Boolean condition is ghost code, then the conditional itself becomes ghost. Note, however, that the typing side-condition (1) will reject conditionals where one part is ghost and another part has some effect, as in

$$\text{if true then } r^{\perp} := 42 \text{ else ghost } ().$$

The rule (T-GHOST) turns any term t into ghost code, with ghost status $\top$, whatever the ghost status of t is, provided that t is pure and terminating. Thus, terms such as `ghost` ($r^{\perp} := 42$) or `ghost` (`fact 3`) are ill-typed, since their evaluation would interfere with the evaluation of regular code.

The side condition $(\beta \geq \beta')$ of the rule (T-ASSIGN) ensures that regular references cannot be assigned ghost code. (Boolean values are ordered as usual, with $\top > \perp$.) Additionally, the rule conclusion ensures that, if the assigned reference is regular ($\beta = \perp$), then ϵ is $\top$; on the contrary, if the assigned reference is ghost ($\beta = \top$), then ϵ is $\perp$, since ghost reference assignments are not part of regular effects.

The most subtle rules are those for local bindings and application. Rule (T-LET-GHOST) states that, whatever the ghost status of a term t_2 is, as long as

t_2 is pure and terminating, we can bind a ghost variable $x^\top$ to t_2 . Similarly, by rule (T-APP-GHOST) a function that expects a ghost parameter can be applied to both ghost and regular values.

Rule (T-LET-REGULAR) is somewhat dual to (T-LET-GHOST): it allows us to bind a regular variable $x^\perp$ to a ghost term. The difference with the previous case is that, now, the ghost status of the `let` expression depends on the ghost status of t_2 : if t_2 is ghost code, then the “contaminated” `let` expression becomes ghost itself. Consequently, if t_2 is ghost, then by the implicit side-condition, as $\epsilon_1 \vee \epsilon_2$ must be equal to $\perp$, both t_1 and t_2 must be pure and terminating. Similarly, rule (T-APP-REGULAR) allows us to pass a ghost value to a function expecting a regular parameter, in which case the application itself becomes ghost. In other words, the goal of rules (T-LET-REGULAR) and (T-APP-REGULAR) is to allow ghost code to use regular code. This was one of our motivations.

It is worth pointing out that there is no sub-typing in our system. That is, in rules for application, the formal parameter and the actual argument must have *exactly* the same type τ_2 . In particular, all latent effects and ghost statuses in function types must be the same. For instance, a function expecting an argument of type $\text{int}^\perp \multimap \text{int}$ cannot be applied to an argument of type $\text{int}^\top \multimap \text{int}$.

Type System of MiniML. Similarly to operational semantics, if we drop all ghost statuses (or, equivalently, if we consider them marked as $\perp$) and get rid of typing rule (T-GHOST), we get a standard typing system with effects for MiniML with simple types. For clarity, we add a subscript m when we write typing judgments for MiniML terms: $\Sigma, \Gamma \vdash_m t : \tau, \epsilon$.

2.4 Type Soundness

The type system of GhostML enjoys the standard soundness property. Any well-typed program either diverges or evaluates to a value. This property is well established in the literature for ML with references [6, 7], and we can easily adapt the proof in our case. Due to lack of space, we only give the main statements.

As usual, we decompose type soundness into *preservation* and *progress* lemmas. First, we define well-typedness of a store with respect to a store typing.

Definition 1. A store μ is well-typed with respect to a store typing Σ , written $\Sigma \vdash \mu$, if $\text{dom}(\mu) \subseteq \text{dom}(\Sigma)$ and $\mu(r^\beta)$ has type $\Sigma(r^\beta)$ for every $r^\beta \in \text{dom}(\mu)$.

With this definition, the *preservation* lemma is stated as follows:

Lemma 1 (Preservation). If $\Sigma, \Gamma \vdash t : \tau, \beta, \epsilon$ and $\Sigma \vdash \mu$, then $t \mid \mu \rightarrow t' \mid \mu'$ implies that $\Sigma, \Gamma \vdash t' : \tau, \beta', \epsilon'$ and $\Sigma \vdash \mu'$, where $\beta \geq \beta'$ and $\epsilon \geq \epsilon'$.

The only difference with respect to the standard statement is that ghost statuses and effect indicators can decrease during evaluation.

Lemma 2 (Progress). If $\Sigma, \emptyset \vdash t : \tau, \beta, \epsilon$, then either t is a value or, for any store μ such that $\Sigma \vdash \mu$, there exists a reduction step $t \mid \mu \rightarrow t' \mid \mu'$.

Additionally, we have the following results for effect-less programs.

Lemma 3 (Store Preservation). *If $\Sigma, \emptyset \vdash t : \tau, \beta, \perp$ and $\Sigma \vdash \mu$, then $t \mid \mu \rightarrow t' \mid \mu'$ implies $\mu_\perp = \mu'_\perp$.*

Lemma 4 (Program Termination). *If $\Sigma, \emptyset \vdash t : \tau, \beta, \perp$ and $\Sigma \vdash \mu$, then evaluation of t in store μ terminates, that is, there is a value v and a store μ' such that $t \mid \mu \rightarrow^* v \mid \mu'$.*

A consequence of the previous lemmas and the side condition (1) is that ghost code does not modify the regular store and is terminating.

3 From GhostML to MiniML

This section describes an erasure operation that turns a GhostML term into a MiniML term. The goal is to show that ghost code can be erased from a regular program without observable difference in the program outcome.

The erasure is written either $\mathcal{E}_\beta(\cdot)$, when parameterized by some ghost status β , and simply $\mathcal{E}(\cdot)$ otherwise. First, we define erasure on types and terms. The main idea is to preserve the structure of regular terms and types, and to replace any ghost code by a value of type `unit`.

Definition 2 (τ -erasure). *Let τ be some GhostML type. The erasure $\mathcal{E}_\beta(\tau)$ of type τ with respect to β is defined by induction on the structure of τ as follows:*

$$\begin{aligned} \mathcal{E}_\top(\tau) &\triangleq \text{unit} \\ \mathcal{E}_\perp(\tau_2^{\beta_2} \multimap \tau_1) &\triangleq \mathcal{E}_{\beta_2}(\tau_2) \multimap \mathcal{E}_\perp(\tau_1) \\ \mathcal{E}_\perp(\kappa) &\triangleq \kappa \end{aligned}$$

In other words, the structure of regular types is preserved and all ghost types are turned into type `unit`. Now we can define erasure on terms.

Definition 3 (t -Erasure). *Let t be such that $\Sigma, \Gamma \vdash t : \tau, \beta, \epsilon$ holds. The erasure $\mathcal{E}_\beta(t)$ is defined by induction on the structure of t as follows:*

$$\begin{aligned} \mathcal{E}_\top(t) &\triangleq () \\ \mathcal{E}_\perp(c) &\triangleq c \\ \mathcal{E}_\perp(x^\perp) &\triangleq x \\ \mathcal{E}_\perp(\lambda x^\beta : \tau. t) &\triangleq \lambda x : \mathcal{E}_\beta(\tau). \mathcal{E}_\perp(t) \\ \mathcal{E}_\perp(\text{rec } f^\perp : \tau_2^{\beta_2} \multimap \tau_1. t) &\triangleq \text{rec } f : \mathcal{E}_\perp(\tau_2^{\beta_2} \multimap \tau_1). \mathcal{E}_\perp(t) \\ \mathcal{E}_\perp(r^\perp := v) &\triangleq r := \mathcal{E}_\perp(v) \\ \mathcal{E}_\perp(!r^\perp) &\triangleq !r \\ \mathcal{E}_\perp(\text{if } v \text{ then } t_1 \text{ else } t_2) &\triangleq \text{if } \mathcal{E}_\perp(v) \text{ then } \mathcal{E}_\perp(t_1) \text{ else } \mathcal{E}_\perp(t_2) \\ \mathcal{E}_\perp(t \ v) &\triangleq \mathcal{E}_\perp(t) \ \mathcal{E}_{\beta'}(v) \quad \text{where } t \text{ has type } \tau_2^{\beta'} \multimap \tau_1 \\ \mathcal{E}_\perp(\text{let } x^{\beta'} = t_2 \text{ in } t_1) &\triangleq \text{let } x = \mathcal{E}_{\beta'}(t_2) \text{ in } \mathcal{E}_\perp(t_1) \end{aligned}$$

Note that ghost variables and ghost references do not occur anymore in $\mathcal{E}_\perp(t)$. Note also that a regular function (recursive or not) with a ghost parameter remains a function, but with an argument of type `unit`. Similarly, a `let` expression that binds a ghost variable inside a regular code remains a `let`, but now binds a variable to `()`. More generally, $\mathcal{E}_\perp(t)$ is a value if and only if t is a value.

Leaving unit values and arguments in the outcome of erasure may seem unnecessary. However, because of latent effects, full erasure of ghost code is not possible. Consider for instance the function

$$\lambda x^\perp : \text{int}. \lambda y^\top : \text{int}. r^\perp := x$$

where r is a regular reference. Then a partial application of this function to a single argument should not trigger the modification of r . Our solution is to keep a second argument y of type `unit`.

3.1 Well-typedness Preservation

We prove that erasure preserves well-typedness of terms. To do so, we first define the erasure of a typing context and of a store typing by a straightforward induction on their size:

Definition 4 (Γ -erasure and Σ -erasure).

$$\begin{aligned} \mathcal{E}(\emptyset) &\triangleq \emptyset & \mathcal{E}(\emptyset) &\triangleq \emptyset \\ \mathcal{E}(\Gamma, x^\top : \tau) &\triangleq \mathcal{E}(\Gamma), x : \text{unit} & \mathcal{E}(\Sigma, r^\top : \kappa) &\triangleq \mathcal{E}(\Sigma) \\ \mathcal{E}(\Gamma, x^\perp : \tau) &\triangleq \mathcal{E}(\Gamma), x : \mathcal{E}_\perp(\tau) & \mathcal{E}(\Sigma, r^\perp : \kappa) &\triangleq \mathcal{E}(\Sigma), r : \kappa \end{aligned}$$

With these definitions, we prove well-typedness preservation under erasure:

Theorem 1 (Well-typedness Preservation). *If $\Sigma, \Gamma \vdash t : \tau$, $\perp, \epsilon$ holds, then $\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash_m \mathcal{E}_\perp(t) : \mathcal{E}_\perp(\tau)$, ϵ holds.*

3.2 Correctness of Erasure

Finally, we prove correctness of erasure, that is, evaluation is preserved by erasure. To turn this into a formal statement, we first define the erasure of a store μ by a straightforward induction on the store size:

Definition 5 (μ -erasure).

$$\begin{aligned} \mathcal{E}(\emptyset) &\triangleq \emptyset \\ \mathcal{E}(\mu \uplus \{r^\top \mapsto c\}) &\triangleq \mathcal{E}(\mu) \\ \mathcal{E}(\mu \uplus \{r^\perp \mapsto c\}) &\triangleq \mathcal{E}(\mu) \uplus \{r \mapsto c\} \end{aligned}$$

Notice that $\mathcal{E}(\mu)$ removes ghost annotations, and thus is not the same that $\mu_\perp$. The correctness of erasure means that, for any evaluation $t \mid \mu \rightarrow^* v \mid \mu'$ in GhostML, we have $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \rightarrow_m^* \mathcal{E}_\perp(v) \mid \mathcal{E}(\mu')$ in MiniML and that, for any

diverging evaluation $t \mid \mu \rightarrow \infty$ in GhostML, we have $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \rightarrow_m \infty$ in MiniML. We prove these two statements using a bisimulation argument. First, we need the substitution lemma below, which states that substitution and erasure commute.

Lemma 5 (Substitution Under Erasure). *Let t be a GhostML term and v a GhostML value such that $\Sigma, \Gamma, x^\beta : \tau \vdash t : \tau_0, \perp, \epsilon$ and $\Sigma, \Gamma \vdash v : \tau, \beta', \perp$, with $\beta \geq \beta'$, hold. Then the following holds:*

$$\mathcal{E}_\perp(t)[x \leftarrow \mathcal{E}_\beta(v)] = \mathcal{E}_\perp(t[x^\beta \leftarrow v]).$$

Note that if $\Sigma \vdash \mu$ then $\mathcal{E}(\Sigma) \vdash_m \mathcal{E}(\mu)$. To prove erasure correctness for terminating programs, we use the following forward simulation argument:

Lemma 6 (Forward Simulation of GhostML). *If $\Sigma, \emptyset \vdash t : \tau, \perp, \epsilon$ and, for some store μ such that $\Sigma \vdash \mu$, we have $t \mid \mu \rightarrow t' \mid \mu'$, then the following holds in MiniML: $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \xrightarrow{0|1}_m \mathcal{E}_\perp(t') \mid \mathcal{E}(\mu')$.*

We are now able to prove the first part of the main theorem:

Theorem 2 (Terminating Evaluation Preservation). *If typing judgment $\Sigma, \emptyset \vdash t : \tau, \perp, \epsilon$ holds and $t \mid \mu \rightarrow^* v \mid \mu'$, for some value v and some store μ such that $\Sigma \vdash \mu$, then $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \xrightarrow{*}_m \mathcal{E}_\perp(v) \mid \mathcal{E}(\mu')$.*

To prove the second part of the erasure correctness (non-termination preservation), we use the following simulation argument.

Lemma 7 (Forward Simulation of MiniML). *If $\Sigma, \emptyset \vdash t : \tau, \perp, \epsilon$ holds, then, for any store μ such that $\Sigma \vdash \mu$, if $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \rightarrow_m q \mid \nu$ for some term q and some store ν , then $t \mid \mu \rightarrow^{\geq 1} t' \mid \mu'$ where $\mathcal{E}_\perp(t') = q$ and $\mathcal{E}(\mu') = \nu$.*

Finally, we establish non-termination preservation:

Theorem 3 (Non-termination Preservation). *If $\Sigma, \emptyset \vdash t : \tau, \perp, \epsilon$ holds and $t \mid \mu \rightarrow \infty$, for some store μ such that $\Sigma \vdash \mu$, then $\mathcal{E}_\perp(t)$ also diverges, that is, $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \rightarrow_m \infty$.*

4 Implementation

Our method to handle ghost code is implemented in the verification tool Why3³. With respect to GhostML, the language and the type system of Why3 have the following extensions:

Type Polymorphism. The type system of Why3 is first-order and features ML-style type polymorphism. Our approach to associate ghost status with variables and expressions, and not with types, makes this extension straightforward.

³ Why3 is freely available from <http://why3.lri.fr/>.

Local References. Another obvious extension of GhostML is the support of non-global references. As long as such a reference cannot be an alias for another one, the type system of GhostML requires practically no changes. In a system where aliases are admitted, the type system and, possibly, the verification condition generator must be adapted to detect modifications made by a ghost code in locations accessible from regular code. In Why3, aliases are tracked statically, and thus non-interference is ensured purely by type checking.

Data Structures with Ghost Fields. Why3 supports algebraic data types (in particular, records), whose fields may be regular or ghost. Pattern matching on such structures requires certain precautions. Any variable bound in the ghost part of a pattern must be ghost. Moreover, pattern matching over a ghost expression that has at least two branches must make the whole expression ghost, whatever the right-hand sides of the branches are, just as in the case of a conditional over a ghost Boolean expression.

That said, ghost code can use the same data types as regular code. A ghost variable may be a record with regular, mutable fields, which can be accessed and modified in ghost code. Similarly, Why3 has a unique type of arrays and admits both regular and ghost arrays.

Exceptions. Adding exceptions is rather straightforward, since in Why3 exceptions are introduced only at the top level. Indeed, it suffices to add a new effect indicator, that is the set of exceptions possibly raised by a program expression. We can use the same exceptions in ghost and regular code, provided that the ghost status of an expression that raises an exception is propagated upwards until the exception is caught.

Provable Termination. For the sake of simplicity, GhostML forbids the use of recursive functions in ghost code. In Why3, the use of recursive functions and loops in ghost code is allowed. The system requires that such constructs are supplied with a “variant” clause, so that verification conditions for termination are generated.

Example. Let us illustrate the use of ghost code in Why3 on a simple example. Fig. 4 contains an implementation of a mutable queue data type, in Baker’s style. A queue is a pair of two immutable singly-linked lists, which serve to amortize push and pop operations. Our implementation additionally stores the pure logical view of the queue as a list, in the third, ghost field of the record. Notice that we use the same `list` type both for regular and ghost data.

We illustrate propagation in function `push` (lines 27–30), where a local variable `v` is used to hold some intermediate value, to be stored later in the ghost field of the structure. Despite the fact that variable `v` is not declared ghost, and the fact that function `append` is a regular function, Why3 infers that `v` is ghost. Indeed, the ghost value `q.view` contaminates the result of `append`. It would therefore generate an error if we tried to store `v` in a non-ghost field of an existing regular structure. Since the expression `append q.view (Cons x Nil)`

```

1 module Queue
2
3   type elt
4
5   type list = Nil | Cons elt list
6
7   let rec append (l1 l2: list) : list
8     variant { l1 }
9   = match l1 with
10    | Nil → l2
11    | Cons x r1 → Cons x (append r1 l2)
12  end
13
14  let rec rev_append (l1 l2: list) : list
15    variant { l1 }
16  = match l1 with
17    | Nil → l2
18    | Cons x r1 → rev_append r1 (Cons x l2)
19  end
20
21  type queue = {
22    mutable front: list;
23    mutable rear: list;
24    ghost mutable view: list;
25  }
26
27  let push (x: elt) (q: queue) : unit
28  = q.rear ← Cons x q.rear;
29    let v = append q.view (Cons x Nil) in
30    q.view ← v
31
32  exception Empty
33
34  let pop (q: queue): elt
35  raises { Empty }
36  = match q.front with
37    | Cons x f →
38      q.front ← f;
39      q.view ← append f (rev_append q.rear Nil);
40      x
41    | Nil →
42      match rev_append q.rear Nil with
43      | Nil →
44        raise Empty
45      | Cons x f →
46        q.front ← f;
47        q.rear ← Nil;
48        q.view ← f;
49        x
50    end
51  end
52 end

```

Fig. 4. Queue implementation in Why3.

is ghost, it must not diverge. Thus Why3 requires function `append` to be terminating. This is ensured by the `variant` clause on line 8. In function `pop` (lines 34–52), the regular function `rev_append` is used both in regular code (line 42) and ghost code (line 39).

The online gallery of verified Why3 programs contains several other examples of use of ghost code⁴, in particular, ghost function parameters and ghost functions to supply automatic induction proofs (also known as lemma functions).

5 Related Work

The idea to use ghost code in a program to ease specification exists since the early days (late sixties) of deductive program verification, when so-called auxiliary variables became a useful technique in the context of concurrent programming. According to Jones [8] and Reynolds [9], the notion of auxiliary variable was first introduced by Lucas in 1968 [10]. Since then, numerous authors have adapted this technique in various domains.

It is worth pointing out that some authors, in particular Kleymann [11] and Reynolds [9], make a clear distinction between non-operational variables used in program annotations and specification-purpose variables that can appear in the program itself. The latter notion has gradually evolved into the wider idea that ghost code can be arbitrary code, provided it does not interfere with regular code. For example, Zhang *et al.* [12] discuss the use of auxiliary code in the context of concurrent program verification. They present a simple WHILE language with parallelism and auxiliary code, and prove that the latter does not interfere with the rest of the program. In their case, non-interference is ensured by the stratified syntax of the language. For instance, loops can contain auxiliary code, but auxiliary code cannot contain loops, which ensures termination. They also define auxiliary code erasure and prove that a program with ghost code has no less behaviors than its regular part. Schmaltz [13] proposes a rigorous description of ghost code for a large fragment of C with parallelism, in the context of the VCC verification tool [2]. VCC includes ghost data types, ghost fields in regular structures, ghost parameters in regular functions, and ghost variables. In particular, ghost code is used to manipulate ownership information. A notable difference w.r.t. our work is that VCC does not perform any kind of inference of ghost code. Another difference is that VCC *assumes* that ghost code terminates, and the presence of constructions such as `ghost(goto 1)` makes it difficult to reason about ghost code termination.

Another example of a modern deductive verification tool implementing ghost code is the program verifier Dafny [1]. In Dafny, “the concept of ghost versus non-ghost declarations is an integral part of the Dafny language: each function, method, variable, and parameter can be declared as either ghost or non-ghost.” [14]. In addition, a class can contain both ghost fields and regular fields. Dafny ensures termination of ghost code. Ghost code can update ghost fields, but

⁴ <http://toccata.lri.fr/gallery/ghost.en.html>

is not allowed to allocate memory or update non-ghost fields. Consequently, ghost code cannot obtain full reuse of libraries that allocate and mutate classes or arrays. However, on the fragment of Dafny’s language corresponding to GhostML, Dafny provides a semantics of ghost code similar to what is presented here.

The property of non-interference of ghost code is a special case of information flow non-interference [15]. Indeed, one can see ghost code as high-security information and regular code as low-security information, and non-interference precisely means that high-security information does not leak into low-security computations. Information flow properties can be checked using a type system [16] and proofs in that domain typically involve a bisimulation technique (though not necessarily through an erasure operation). Notice that applying an information flow type system to solve our problem is not straightforward, since termination of ghost code is a crucial requirement. For instance, the type system described by Simonet and Pottier [17] simply assumes termination of secret code. To the best of our knowledge, this connection between information flow and ghost code has not been made before, and mainstream deductive verification tools employ syntactical criteria of non-interference instead of type-based ones. In this paper, we develop such a type-based approach, specifically tailored for program verification.

6 Conclusion and Perspectives

In this paper, we described an ML-like language with ghost code. Non-interference between ghost code and regular code is ensured using a type system with effects. We formally proved the soundness of this type system, that is, ghost code can be erased without observable difference. Our type system results in a highly expressive language, where the same data types and functions can be reused in both ghost and regular code.

We see two primary directions of future work on ghost code and Why3. First, ghost code, especially ghost fields, plays an important role in program refinement. Indeed, ghost fields that give sufficient information to specify a data type are naturally shared between the interface and the implementation of this data type. In this way, the glue invariant becomes nothing more than the data type invariant linking regular and ghost fields together. Our intention is to design and implement in Why3 a module system with refinement that makes extensive use of ghost code and data. Second, since ghost code does not have to be executable, it should be possible to use in ghost code various constructs which, up to now, may only appear in specifications, such as quantifiers, inductive predicates, non-deterministic choice, or infinitely parallel computations (cf. the aggregate `forall` statement in Dafny).

Acknowledgments. We are grateful to Sylvain Conchon, Rustan Leino, and François Pottier for comments and discussions regarding earlier versions of this paper.

References

1. Leino, K.R.M.: Dafny: An Automatic Program Verifier for Functional Correctness. In Springer, ed.: LPAR-16. Volume 6355. (2010) 348–370
2. Cohen, E., Dahlweid, M., Hillebrand, M., Leinenbach, D., Moskal, M., Santen, T., Schulte, W., Tobies, S.: VCC: A practical system for verifying concurrent C. In: Theorem Proving in Higher Order Logics (TPHOLs). Volume 5674 of Lecture Notes in Computer Science., Springer (2009)
3. Jacobs, B., Piessens, F.: The VeriFast program verifier. CW Reports CW520, Department of Computer Science, K.U.Leuven (August 2008)
4. Filliâtre, J.C., Paskevich, A.: Why3 — where programs meet provers. In Felleisen, M., Gardner, P., eds.: Proceedings of the 22nd European Symposium on Programming. Volume 7792 of Lecture Notes in Computer Science., Springer (March 2013) 125–128
5. Flanagan, C., Sabry, A., Duba, B.F., Felleisen, M.: The essence of compiling with continuations. SIGPLAN Not. **28**(6) (June 1993) 237–247
6. Wright, A.K., Felleisen, M.: A syntactic approach to type soundness. Information and Computation **115** (1992) 38–94
7. Pierce, B.C.: Types and Programming Languages. MIT Press (2002)
8. Jones, C.B., Roscoe, A., Wood, K.R.: Reflections on the Work of C.A.R. Hoare. 1st edn. Springer Publishing Company, Incorporated (2010)
9. Reynolds, J.C.: The craft of programming. Prentice Hall International series in computer science. Prentice Hall (1981)
10. Lucas, P.: Two constructive realizations of the block concept and their equivalence. Technical Report 25.085, IBM Laboratory, Vienna (June 1968)
11. Kleymann, T.: Hoare logic and auxiliary variables. Formal Asp. Comput. **11**(5) (1999) 541–566
12. Zhang, Z., Feng, X., Fu, M., Shao, Z., Li, Y.: A structural approach to prophecy variables. In Agrawal, M., Cooper, S., Li, A., eds.: 9th annual conference on Theory and Applications of Models of Computation (TAMC). Volume 7287 of Lecture Notes in Computer Science., Springer Berlin Heidelberg (2012) 61–71
13. Schmaltz, S.: Towards the Pervasive Formal Verification of Multi-Core Operating Systems and Hypervisors Implemented in C. PhD thesis, Saarland University, Saarbrücken (2013)
14. Leino, K.R.M., Moskal, M.: Co-induction simply: Automatic co-inductive proofs in a program verifier. In Jones, C., Pihlajasaari, P., Sun, J., eds.: FM 2014: Formal Methods. Volume 8442 of Lecture Notes in Computer Science., Springer (May 2014) 382–398
15. Denning, D.E., Denning, P.J.: Certification of programs for secure information flow. Communications of the ACM **20**(2) (July 1977) 504–513
16. Pottier, F., Conchon, S.: Information flow inference for free. In: Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP’00), Montréal, Canada (September 2000) 46–57
17. Pottier, F., Simonet, V.: Information flow inference for ML. ACM Transactions on Programming Languages and Systems **25**(1) (January 2003) 117–158 ©ACM.