



HAL
open science

The Spirit of Ghost Code

Jean-Christophe Filliâtre, Léon Gondelman, Andrei Paskevich

► **To cite this version:**

Jean-Christophe Filliâtre, Léon Gondelman, Andrei Paskevich. The Spirit of Ghost Code. 2013. hal-00873187v1

HAL Id: hal-00873187

<https://inria.hal.science/hal-00873187v1>

Preprint submitted on 16 Oct 2013 (v1), last revised 14 May 2014 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

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 any 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 compute verification conditions, for instance using a weakest preconditions 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 computation,

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

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].

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.

2 GhostML

We introduce GhostML, a mini ML-like language with ghost code. It features global references, 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, etc. Each variable is tagged with a

$t ::=$	TERMS	$v ::=$	VALUES
v	value	c	constant
$(t v)$	application	x^β	variable
$\text{let } x^\beta = t \text{ in } t$	local binding	$\lambda x^\beta : \tau. t$	anonymous function
$\text{if } v \text{ then } t \text{ else } t$	conditional	$\text{rec } x^\beta : \tau^\beta \stackrel{(\theta, \rho)}{\implies} \tau. t$	recursive function
$r^\beta := v$	assignment	$()$	unit
$!r^\beta$	dereference	$\dots, -1, 0, 1, \dots$	integers
$\text{ghost } t$	ghost code	$\text{true}, \text{false}$	Boolean
$\tau ::=$	TYPES	$+, \vee, =, \dots$	operators
κ	built-in type	$\theta \in \{\perp, \top\}$	EFFECTS
$\text{ref } \kappa$	reference type	$\rho \in \{\perp, \top\}$	assignment
$\tau^\beta \stackrel{(\theta, \rho)}{\implies} \tau$	functional type	$\beta \in \{\perp, \top\}$	recursive call
$\kappa ::=$	BUILT-IN TYPES		GHOST STATUS
$\text{int} \mid \text{bool} \mid \text{unit}$	built-in types		

Fig. 1. Syntax.

ghost status β , which is \top for ghost variables and \perp for regular ones. Terms can assign or access global references. We assume that a finite set of global references with distinct names is given. Similar to local variables, references are tagged with a ghost status. Note that each *occurrence* of a reference or a variable is explicitly annotated with its ghost status. Formal parameters of functions are also introduced with a ghost status. For instance, the following term introduces a local function upd^\top , that updates a ghost reference g^\top with its parameter, and applies it to the contents of a regular reference r^\perp :

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

The keyword `ghost` turns a term t into ghost code. For instance, `ghost 42` is a ghost term. This way, the same constant 42 can be used in both regular and ghost code.

Note that compound terms obey to 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 $\text{let } x^\perp = (t_2 v) \text{ in } (t_1 x^\perp)$.

Types are either primitive data-types (`int`, `bool`, `unit`), reference types, or function types. A function type is an arrow $\tau_2^\beta \stackrel{(\theta, \rho)}{\implies} \tau_1$ where β stands for the function argument's ghost status, and (θ, ρ) is the *latent* effect of the function. θ stands for the presence of non-ghost reference assignment, and ρ stands for the presence of 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 μ_{\perp} , is the restriction of μ to regular references. Rules indicate the store μ only when relevant.

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

Fig. 2. Semantics.

A reduction step can take place directly *on the top* of a term t . Such step is called a head reduction and is denoted $t \mid \mu \xrightarrow{e}_g t \mid \mu'$.

The rule (E-GHOST) expresses that, from the semantics point of view, there is no difference between regular and ghost code. Other head reduction rules are standard. For instance, rules (E-CONST- λ) and (E-CONST- 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 $\lambda x^\perp : \kappa. c_1 \dots c_m$, 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. Note also that the rule (E-APP-REC) only reduces recursive functions whose body is a function. Said otherwise, only functions can be defined recursively.

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_g^* denotes the reflexive, transitive closure of \rightarrow_g . We say that a closed term t evaluates to v in a store μ if there exists μ' such that $t \mid \mu \rightarrow_g^* v \mid \mu'$. 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:

$$\frac{t \mid \mu \xrightarrow{1}_g t' \mid \mu' \quad t' \mid \mu' \rightarrow_g \infty}{t \mid \mu \rightarrow_g \infty} \text{(E-DIV)}$$

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 SOS semantics for MiniML.

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 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

$$\Sigma, \Gamma \vdash_g t : \tau, \beta, (\theta, \rho).$$

Here, Σ is a store typing, that binds references to types, and Γ is a typing environment, that binds variables to types. Boolean indicators β , θ , and ρ indicate, respectively, the ghost status of t , whether some *regular* reference is possibly assigned by t , and whether t is possibly non-terminating.

Typing rules are given in Fig. 3. In each rule, whose conclusion is a judgment $\Sigma, \Gamma \vdash_g t : \tau, \beta, (\theta, \rho)$, we add the implicit extra side-condition

$$(\beta = \top) \Rightarrow (\theta = \perp \wedge \rho = \perp) \tag{1}$$

$\frac{\text{Typeof}(c) = \tau}{\Sigma, \Gamma \vdash_{\mathbf{g}} c : \tau, \perp, (\perp, \perp)}$	(T _g -CONST)
$\frac{x^\beta : \tau \in \Gamma}{\Sigma, \Gamma \vdash_{\mathbf{g}} x^\beta : \tau, \beta, (\perp, \perp)}$	(T _g -VAR)
$\frac{\Sigma, \Gamma, x^\beta : \tau \vdash_{\mathbf{g}} t : \tau_0, \beta_0, (\theta, \rho)}{\Sigma, \Gamma \vdash_{\mathbf{g}} \lambda x^\beta : \tau. t : \tau^\beta \xrightarrow{(\theta, \rho)} \tau_0, \beta_0, (\perp, \perp)}$	(T _g -λ)
$\frac{\Sigma, \Gamma, f^\perp : \tau \vdash_{\mathbf{g}} \lambda x^\beta : \tau_2. t : \tau_2^\beta \xrightarrow{(\theta, \rho)} \tau_1, \perp, (\perp, \perp) \quad \text{where } \tau \triangleq \tau_2^\beta \xrightarrow{(\theta, \top)} \tau_1}{\Sigma, \Gamma \vdash_{\mathbf{g}} \text{rec } f^\perp : \tau. \lambda x^\beta : \tau_2. t : \tau, \perp, (\perp, \perp)}$	(T _g -REC)
$\frac{\Sigma, \Gamma \vdash_{\mathbf{g}} v : \text{bool}, \beta_0, (\perp, \perp) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} t_1 : \tau, \beta_1, (\theta_1, \rho_1) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} t_2 : \tau, \beta_2, (\theta_2, \rho_2)}{\Sigma, \Gamma \vdash_{\mathbf{g}} \text{if } v \text{ then } t_1 \text{ else } t_2 : \tau, \beta_0 \vee \beta_1 \vee \beta_2, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$	(T _g -IF)
$\frac{\Sigma, \Gamma, x^\perp : \tau_2 \vdash_{\mathbf{g}} t_1 : \tau_1, \beta_1, (\theta_1, \rho_1) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} t_2 : \tau_2, \beta_2, (\theta_2, \rho_2)}{\Sigma, \Gamma \vdash_{\mathbf{g}} \text{let } x^\perp = t_2 \text{ in } t_1 : \tau_1, \beta_1 \vee \beta_2, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$	(T _g -LET-REGULAR)
$\frac{\Sigma, \Gamma, x^\top : \tau_2 \vdash_{\mathbf{g}} t_1 : \tau_1, \beta_1, (\theta_1, \rho_1) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} t_2 : \tau_2, \beta_2, (\perp, \perp)}{\Sigma, \Gamma \vdash_{\mathbf{g}} \text{let } x^\top = t_2 \text{ in } t_1 : \tau_1, \beta_1, (\theta_1, \rho_1)}$	(T _g -LET-GHOST)
$\frac{\Sigma, \Gamma \vdash_{\mathbf{g}} t : \tau_2^\perp \xrightarrow{(\theta_1, \rho_1)} \tau_1, \beta_1, (\theta_2, \rho_2) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} v : \tau_2, \beta_2, (\perp, \perp)}{\Sigma, \Gamma \vdash_{\mathbf{g}} (t \ v) : \tau_1, \beta_1 \vee \beta_2, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$	(T _g -APP-REGULAR)
$\frac{\Sigma, \Gamma \vdash_{\mathbf{g}} t : \tau_2^\top \xrightarrow{(\theta_1, \rho_1)} \tau_1, \beta_1, (\theta_2, \rho_2) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} v : \tau_2, \beta_2, (\perp, \perp)}{\Sigma, \Gamma \vdash_{\mathbf{g}} (t \ v) : \tau_1, \beta_1, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$	(T _g -APP-GHOST)
$\frac{r^\beta : \text{ref } \kappa \in \Sigma}{\Sigma, \Gamma \vdash_{\mathbf{g}} !r^\beta : \kappa, \beta, (\perp, \perp)}$	(T _g -DEREF)
$\frac{\Sigma, \Gamma \vdash_{\mathbf{g}} v : \kappa, \beta', (\perp, \perp) \quad r^\beta : \text{ref } \kappa \in \Sigma \quad \beta' \Rightarrow \beta}{\Sigma, \Gamma \vdash_{\mathbf{g}} r^\beta := v : \text{unit}, \beta, (\neg\beta, \perp)}$	(T _g -ASSIGN)
$\frac{\Sigma, \Gamma \vdash_{\mathbf{g}} t : \tau, \beta, (\perp, \perp)}{\Sigma, \Gamma \vdash_{\mathbf{g}} \text{ghost } t : \tau, \top, (\perp, \perp)}$	(T _g -GHOST)

Fig. 3. Typing rules.

to account for non-interference. Said otherwise, whenever t is ghost code, then it must terminate and must not assign any regular reference.

Let us explain some rules in details. The rule (T_g -CONST) states that any constant c is regular code ($\beta = \perp$) yet is pure and terminating ($\theta = \perp, \rho = \perp$). The type of each constant is given by some oracle function $\text{Typeof}(c)$. For instance, $\text{Typeof}(+) = \text{int}^{\perp (\frac{\perp+\perp}{\perp})} \text{int}^{\perp (\frac{\perp+\perp}{\perp})} \text{int}$.

Recursive functions are typed as follows. For simplicity, we assume that whenever a recursive function is used, we may have non-termination. (In Section 4, we will discuss more a elaborate solution.) To do that, we enforce the latent effect ρ of any recursive function to be \top . Consequently, we declare any recursive function as non-ghost. Note that the body of a recursive function could be terminating, as in $\text{rec } f^{\perp} : \tau. \lambda x^{\perp} : \text{int}. x^{\perp}$ yet we give it the type $\tau = \text{int}^{\perp (\frac{\perp+\top}{\perp})} \text{int}$ with a non-termination effect.

The rule (T_g -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

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

The rule (T_g -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 $\text{ghost } (r^{\perp} := 42)$ or $\text{ghost } (\text{fact } 3)$ are ill-typed, since their evaluation would interfere with the the evaluation of regular code.

The side-condition of the rule (T_g -ASSIGN) checks that regular references cannot be assigned ghost code. Additionally, the rule conclusion ensures that, if the assigned reference is regular ($\beta = \perp$), then θ is \top ; on the contrary, if the assigned reference is ghost ($\beta = \top$), then θ is \perp , since ghost reference assignments are not part of effects.

The most subtle rules are those for local bindings and application. Rule (T_g -LET-GHOST) states that, whatever the ghost status of a term t_2 is, as long as a term t_1 is pure and terminating, we can bind a ghost variable x^{\top} to t_2 . Similarly, Rule (T_g -APP-GHOST) allows a function with a ghost parameter to be passed a regular value.

Rule (T_g -LET-REGULAR) is somehow dual to (T_g -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 $\theta_1 \vee \theta_2$ and $\rho_1 \vee \rho_2$ must both be equal to \perp , both t_1 and t_2 must be pure and terminating. Similarly, rule (T_g -APP-REGULAR) allows us to pass a ghost value to a function with a regular parameter, in which case the application itself becomes ghost. In other words, the goal of rules (T_g -LET-REGULAR) and (T_g -LET-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 $\text{int}^{\perp (\underline{\theta}, \underline{\rho})}$ cannot be applied to an argument of type $\text{int}^{\top (\underline{\theta}, \underline{\rho})}$, though this would be safe.

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_g-GHOST), we get a standard typing system with effects for MiniML with simple types.

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. To do so, we first define well-typedness of a store with respect to a store typing.

Definition 1. A store μ is said to be well typed with respect to a store typing Σ , written $\Sigma \vdash_g \mu$, 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 $\Sigma, \Gamma \vdash_g p : \tau, \beta, (\theta, \rho)$ and $\Sigma \vdash_g \mu$, then $p \mid_{\mu} \rightarrow_g p' \mid_{\mu'}$ implies $\Sigma, \Gamma \vdash_g p' : \tau, \beta', (\theta', \rho')$ and $\Sigma \vdash_g \mu'$, where $\beta \geq \beta'$, $\theta \geq \theta'$, and $\rho \geq \rho'$.

The only difference with respect to the standard statement is that ghost statuses and effect indicators cannot increase during evaluation. (Boolean values are ordered as usual, with $\perp \leq \top$.)

Lemma 2 (Progress). If $\Sigma, \emptyset \vdash_g t : \tau, \beta, (\theta, \rho)$, then either t is a value or, for any store μ such that $\Sigma \vdash_g \mu$, there exists a reduction step $t \mid_{\mu} \rightarrow_g t' \mid_{\mu'}$.

Additionally we have the following results.

Lemma 3 (Store Preservation). If $\Sigma, \emptyset \vdash_g t : \tau, \beta, (\perp, \rho)$ and, for some μ , we have $\Sigma \vdash_g \mu$ and $t \mid_{\mu} \rightarrow_g t' \mid_{\mu'}$, then $\mu_{\perp} = \mu'_{\perp}$.

Lemma 4 (Program Termination). If $\Sigma, \emptyset \vdash_g t : \tau, \beta, (\theta, \perp)$, then evaluation of t terminates. That is, for any store μ such that $\Sigma \vdash_g \mu$, there is a value v and a store μ' such that $t \mid_{\mu} \rightarrow_g^* v \mid_{\mu'}$.

A consequence of the previous two lemmas is that ghost code does not modify the regular store and is terminating.

Corollary 1 (Ghost Code Behavior). *If $\Sigma, \emptyset \vdash_{\mathbf{g}} t : \tau, \top, (\perp, \perp)$, then, for any state $t \mid \mu$ such that $\Sigma \vdash_{\mathbf{g}} \mu$, there exists a value v and a store μ' such that $t \mid \mu \rightarrow_{\mathbf{g}}^* v \mid \mu'$ and $\mu|_{\perp} = \mu'|_{\perp}$.*

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 the ghost code can be erased from a regular program without any observable difference in the program outcome.

The erasure is written either $\mathcal{E}_{\beta}(\cdot)$, when parameterized by some ghost status β , or 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:*

$$\begin{aligned} \mathcal{E}_{\top}(\tau) &= \text{unit} \\ \mathcal{E}_{\perp}(\tau_2^{\beta_2} \stackrel{(\theta, \rho)}{\Longrightarrow} \tau_1) &= \mathcal{E}_{\beta_2}(\tau_2) \stackrel{(\theta, \rho)}{\Longrightarrow} \mathcal{E}_{\perp}(\tau_1) \\ \mathcal{E}_{\perp}(\kappa) &= \kappa \\ \mathcal{E}_{\perp}(\text{ref } \kappa) &= \text{ref } \kappa \end{aligned}$$

Said otherwise, 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 $\Sigma, \Gamma \vdash_{\mathbf{g}} t : \tau, \beta, (\theta, \rho)$ holds. The erasure $\mathcal{E}_{\beta}(t)$ is defined by induction on the structure of t as follows:*

$$\begin{aligned} \mathcal{E}_{\top}(t) &= () \\ \mathcal{E}_{\perp}(c) &= c \\ \mathcal{E}_{\perp}(x^{\perp}) &= x \\ \mathcal{E}_{\perp}(\lambda x^{\beta} : \tau. t) &= \lambda x : \mathcal{E}_{\beta}(\tau). \mathcal{E}_{\perp}(t) \\ \mathcal{E}_{\perp}(\text{rec } f^{\perp} : \tau_2^{\beta_2} \stackrel{(\theta, \top)}{\Longrightarrow} \tau_1. t) &= \text{rec } f : \mathcal{E}_{\perp}(\tau_2^{\beta_2} \stackrel{(\theta, \top)}{\Longrightarrow} \tau_1). \mathcal{E}_{\perp}(t) \\ \mathcal{E}_{\perp}(r^{\perp} := v) &= r := \mathcal{E}_{\perp}(v) \\ \mathcal{E}_{\perp}(!r^{\perp}) &= !r \\ \mathcal{E}_{\perp}(\text{if } v \text{ then } t_1 \text{ else } t_2) &= \text{if } \mathcal{E}_{\perp}(v) \text{ then } \mathcal{E}_{\perp}(t_1) \text{ else } \mathcal{E}_{\perp}(t_2) \\ \mathcal{E}_{\perp}(t \ v) &= \mathcal{E}_{\perp}(t) \ \mathcal{E}_{\beta'}(v) \quad \text{where } t \text{ has type } \tau_2^{\beta'} \stackrel{(\theta_{\perp}, \rho_1)}{\Longrightarrow} \tau_1 \\ \mathcal{E}_{\perp}(\text{let } x^{\beta'} = t_2 \text{ in } t_1) &= \text{let } x = \mathcal{E}_{\beta'}(t_2) \text{ in } \mathcal{E}_{\perp}(t_1) \end{aligned}$$

Note that a ghost variable x^{\top} or a ghost reference r^{\top} does 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 it 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 intrusive. Indeed, after instrumenting a MiniML program with ghost code (for the purpose of verification) and then erasing it, we will not recover the original program. However, the program after erasure is semantically equivalent to the original one, with only extra administrative reduction steps (a term coined by Plotkin [8]) that do not impact the program complexity and can be eliminated at compile time. These reductions facilitate the proofs of forthcoming theorems.

3.1 Well-typedness Preservation

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

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

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

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

Theorem 1 (Well-typedness Preservation). *If $\Sigma, \Gamma \vdash_{\mathbf{g}} t : \tau, \perp, (\theta, \rho)$ holds, then $\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(t) : \mathcal{E}_\perp(\tau), (\theta, \rho)$ holds.*

Proof. By induction on the typing derivation, with case analysis on the last applied rule.

Case (T_g-λ).
$$\frac{\Sigma, \Gamma, x^\beta : \tau \vdash_{\mathbf{g}} t : \tau_0, \perp, (\theta, \rho)}{\Sigma, \Gamma \vdash_{\mathbf{g}} \lambda x^\beta : \tau. t : \tau^\beta \xrightarrow{(\theta, \rho)} \tau_0, \perp, (\perp, \perp)}$$

By induction hypothesis, we have $\mathcal{E}(\Sigma), \mathcal{E}(\Gamma), x : \mathcal{E}_\beta(\tau) \vdash \mathcal{E}_\perp(t) : \mathcal{E}_\perp(\tau_0), (\theta, \rho)$. Therefore, we conclude that

$$\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \lambda x : \mathcal{E}_\beta(\tau). \mathcal{E}_\perp(t) : \mathcal{E}_\beta(\tau) \xrightarrow{(\theta, \rho)} \mathcal{E}_\perp(\tau_0), (\perp, \perp)$$

Case (T_g-APP-REGULAR).

$$\frac{\Sigma, \Gamma \vdash_{\mathbf{g}} t : \tau_2^\perp \xrightarrow{(\theta_1, \rho_1)} \tau_1, \beta_1, (\theta_2, \rho_2) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} v : \tau_2, \beta_2, (\perp, \perp)}{\Sigma, \Gamma \vdash_{\mathbf{g}} (t v) : \tau_1, \beta_1 \vee \beta_2, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$$

By hypothesis $\beta_1 \vee \beta_2 = \perp$. By induction hypotheses on the rule premises,

$$\begin{aligned} \mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(t) &: \mathcal{E}_\perp(\tau_2) \xrightarrow{(\theta_1, \rho_1)} \mathcal{E}_\perp(\tau_1), (\theta_2, \rho_2) \\ \mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(v) &: \mathcal{E}_\perp(\tau_2), (\perp, \perp). \end{aligned}$$

Thus, $\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(t) \mathcal{E}_\perp(v) : \mathcal{E}_\perp(\tau_1), (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)$

Case (T_g-APP-GHOST).

$$\frac{\Sigma, \Gamma \vdash_{\mathbf{g}} t : \tau_2^\top \xrightarrow{(\theta_1, \rho_1)} \tau_1, \beta_1, (\theta_2, \rho_2) \quad \Sigma, \Gamma \vdash_{\mathbf{g}} v : \tau_2, \beta_2, (\perp, \perp)}{\Sigma, \Gamma \vdash_{\mathbf{g}} (t v) : \tau_1, \beta_1, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$$

By hypothesis $\beta_1 = \perp$. Thus, $\mathcal{E}_\perp(t v) = \mathcal{E}_\perp(t) ()$. By induction hypothesis,

$$\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(t) : \text{unit} \stackrel{(\theta_1, \rho_1)}{\Rightarrow} \mathcal{E}_\perp(\tau_1), (\theta_2, \rho_2)$$

Since $\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash () : \text{unit}, (\perp, \perp)$, we conclude that

$$\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(t) () : \mathcal{E}_\perp(\tau_1), (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)$$

Case (T_g-Assign):
$$\frac{\Sigma, \Gamma \vdash_g v : \kappa, \beta', (\perp, \perp) \quad r^\beta : \text{ref } \kappa \in \Sigma \quad \beta' \Rightarrow \beta}{\Sigma, \Gamma \vdash_g r^\beta := v : \text{unit}, \beta, (\neg\beta, \perp)}$$

By hypothesis $\beta = \perp$, and by premise side-condition $\beta' \Rightarrow \beta$, so $\beta' = \perp$ as well. Also, $r^\beta : \text{ref } \kappa \in \Sigma$, so $r : \text{ref } \kappa \in \mathcal{E}(\Sigma)$. Thus, by induction hypothesis on the rule premise, we have $\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(v) : \kappa, (\perp, \perp)$. Therefore:

$$\frac{\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash \mathcal{E}_\perp(v) : \kappa, (\perp, \perp) \quad r : \text{ref } \kappa \in \mathcal{E}(\Sigma)}{\mathcal{E}(\Sigma), \mathcal{E}(\Gamma) \vdash r := \mathcal{E}_\perp(v) : \text{unit}, (\top, \perp)}$$

Other cases are easily proved in a similar way. ■

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).

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

The correctness of erasure means that, for any evaluation $t \mid \mu \xrightarrow*_g v \mid \mu'$ in GhostML, we have $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \xrightarrow*_g \mathcal{E}_\perp(v) \mid \mathcal{E}(\mu')$ in MiniML and that, for any diverging evaluation $t \mid \mu \xrightarrow*_g \infty$ in GhostML, we have $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \rightarrow \infty$ 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 some GhostML term and v some GhostML value such that $\Sigma, \Gamma, x^\beta : \tau \vdash_g t : \tau_0, \perp, (\theta, \rho)$ and $\Sigma, \Gamma \vdash_g v : \tau, \beta', (\perp, \perp)$ where $\beta \geq \beta'$, hold. Then the following holds:*

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

Proof. By straightforward induction on the structure of t . ■

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

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

Proof. By induction on the derivation \rightarrow_g .

CASE (E-HEAD). By case analysis on the head reduction $\xrightarrow{\xi}_g$. We only give the proof for (E-APP- λ); other cases are proved in a similar way.

SUB-CASE (E-APP- λ). $(\lambda x^\beta : \tau. t) v \xrightarrow{\xi}_g t[x^\beta \leftarrow v]$.

If $\beta = \top$, then

$$\frac{\Sigma, \Gamma \vdash_g \lambda x^\top : \tau_2. t : \tau_2^\top \xrightarrow{(\theta_1, \rho_1)} \tau_1, \perp, (\theta_2, \rho_2) \quad \Sigma, \Gamma \vdash_g v : \tau_2, \beta_2, (\perp, \perp)}{\Sigma, \Gamma \vdash_g (\lambda x^\top : \tau_2. t) v : \tau_1, \perp, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$$

Therefore, by substitution under erasure (Lemma 5), we get

$$\mathcal{E}_\perp((\lambda x^\top : \tau. t) v) = (\lambda x : \text{unit}. \mathcal{E}_\perp(t)) () \xrightarrow{\xi} \mathcal{E}_\perp(t)[x \leftarrow ()] = \mathcal{E}_\perp(t[x^\top \leftarrow v]).$$

If $\beta = \perp$, then we have

$$\frac{\Sigma, \Gamma \vdash_g \lambda x^\perp : \tau_2. t : \tau_2^\perp \xrightarrow{(\theta_1, \rho_1)} \tau_1, \perp, (\theta_2, \rho_2) \quad \Sigma, \Gamma \vdash_g v : \tau_2, \perp, (\perp, \perp)}{\Sigma, \Gamma \vdash_g (\lambda x^\perp : \tau_2. t) v : \tau_1, \perp, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$$

Again, by substitution under erasure (Lemma 5), we get

$$\mathcal{E}_\perp((\lambda x^\perp : \tau. t) v) \xrightarrow{\xi} \mathcal{E}_\perp(t)[x \leftarrow \mathcal{E}_\perp(v)] = \mathcal{E}_\perp(t[x^\perp \leftarrow v]).$$

CASE (E-CONTEXT-LET). $\frac{t_2 \mid \mu \rightarrow_g t'_2 \mid \mu'}{\text{let } x^\beta = t_2 \text{ in } t_1 \mid \mu \rightarrow_g \text{let } x^\beta = t'_2 \text{ in } t_1 \mid \mu'}$.

If $\beta = \top$, then $\frac{\Sigma, \Gamma, x^\top : \tau_2 \vdash_g t_1 : \tau_1, \perp, (\theta_1, \rho_1) \quad \Sigma, \Gamma \vdash_g t_2 : \tau_2, \beta_2, (\perp, \perp)}{\Sigma, \Gamma \vdash_g \text{let } x^\top = t_2 \text{ in } t_1 : \tau_1, \perp, (\theta_1, \rho_1)}$

By store preservation (Lemma 3), $\mu|_\perp = \mu'|_\perp$, so $\mathcal{E}(\mu) = \mathcal{E}(\mu')$. Therefore,

$$\mathcal{E}_\perp(\text{let } x^\top = t_2 \text{ in } t_1) \mid \mathcal{E}(\mu) \xrightarrow{0}_g \mathcal{E}_\perp(\text{let } x^\top = t'_2 \text{ in } t_1) \mid \mathcal{E}(\mu').$$

If $\beta = \perp$, then $\frac{\Sigma, \Gamma, x^\perp : \tau_2 \vdash_g t_1 : \tau_1, \perp, (\theta_1, \rho_1) \quad \Sigma, \Gamma \vdash_g t_2 : \tau_2, \perp, (\theta_2, \rho_2)}{\Sigma, \Gamma \vdash_g \text{let } x^\perp = t_2 \text{ in } t_1 : \tau_1, \perp, (\theta_1 \vee \theta_2, \rho_1 \vee \rho_2)}$

By induction hypothesis on sub-derivation $t_2 \mid \mu \rightarrow_g t'_2 \mid \mu'$, we get

$$\mathcal{E}_\perp(t_2) \mid \mathcal{E}_\perp(\mu) \xrightarrow{0|1} \mathcal{E}_\perp(t'_2) \mid \mathcal{E}_\perp(\mu').$$

The result trivially holds when this reduction has 0 step. Otherwise, we have

$$\frac{\mathcal{E}_\perp(t_2) \mid \mathcal{E}_\perp(\mu) \rightarrow \mathcal{E}_\perp(t'_2) \mid \mathcal{E}_\perp(\mu')}{\text{let } x = \mathcal{E}_\perp(t_2) \text{ in } \mathcal{E}_\perp(t_1) \mid \mathcal{E}_\perp(\mu) \rightarrow \text{let } x = \mathcal{E}_\perp(t'_2) \text{ in } \mathcal{E}_\perp(t_1) \mid \mathcal{E}_\perp(\mu')}$$

The remaining case (E-CONTEXT-APP) is similar. ■

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

Theorem 2 (Terminating Evaluation Preservation).

If $\Sigma, \emptyset \vdash_g t : \tau, \perp$, (θ, ρ) holds and $t \mid_{\mu} \rightarrow_g^* v \mid_{\mu'}$, for some value v and some store μ such that $\Sigma \vdash_g \mu$, then $\mathcal{E}_{\perp}(t) \mid_{\mathcal{E}_{\perp}(\mu)} \rightarrow^* \mathcal{E}_{\perp}(v) \mid_{\mathcal{E}_{\perp}(\mu')}$.

Proof. By induction on the length of the evaluation $t \mid_{\mu} \rightarrow_g^* v \mid_{\mu'}$. If $t \mid_{\mu} \rightarrow_g^0 v \mid_{\mu'}$, then the result trivially holds since $t = v$ and $\mu' = \mu$. Now, assume that $t \mid_{\mu} \rightarrow_g^1 t'' \mid_{\mu''} \rightarrow_g^n v \mid_{\mu'}$ for some intermediate store μ'' . By preservation (Lemma 1), $\Sigma, \Gamma \vdash_g t'' : \tau, \perp$, (θ'', ρ'') and $\Sigma \vdash_g \mu''$ for some θ'' and ρ'' , such that $\theta' \geq \theta''$ and $\rho' \geq \rho''$. By induction hypothesis on t'' , $\mathcal{E}_{\perp}(t'') \mid_{\mathcal{E}_{\perp}(\mu'')} \rightarrow^* \mathcal{E}_{\perp}(v) \mid_{\mathcal{E}_{\perp}(\mu')}$. By forward simulation (Lemma 6), $\mathcal{E}_{\perp}(t) \mid_{\mathcal{E}_{\perp}(\mu)} \rightarrow^{0|1} \mathcal{E}_{\perp}(t'') \mid_{\mathcal{E}_{\perp}(\mu'')}$. Putting pieces together we conclude $\mathcal{E}_{\perp}(t) \mid_{\mathcal{E}_{\perp}(\mu)} \rightarrow^{0|1} \mathcal{E}_{\perp}(t'') \mid_{\mathcal{E}_{\perp}(\mu'')} \rightarrow^* \mathcal{E}_{\perp}(v) \mid_{\mathcal{E}_{\perp}(\mu')}$. ■

To prove the second part of the main theorem (non-termination preservation), we use the following backward simulation argument.

Lemma 7 (Backward Simulation). If $\Sigma, \emptyset \vdash_g t : \tau, \perp$, (θ, ρ) holds, then, for any store μ such that $\Sigma \vdash_g \mu$, if $\mathcal{E}_{\perp}(t) \mid_{\mathcal{E}(\mu)} \rightarrow q \mid_{\nu}$ for some term q and some state ν , then $t \mid_{\mu} \rightarrow_g^{\geq 1} t' \mid_{\mu'}$ where $\mathcal{E}_{\perp}(t') = q$ and $\mathcal{E}(\mu') = \nu$.

Proof. By induction on the term t .

Case 1: t is $\text{let } x^{\top} = t_2 \text{ in } t_1$. If t_2 is not a value, then by GhostML progress (Lemma 2), it can be reduced. By hypothesis, t is well-typed, with a typing judgment for t_2 being $\Sigma, \emptyset \vdash_g t_2 : \tau, \beta, (\perp, \perp)$. By program termination (Lemma 4), we have that $t_2 \mid_{\mu} \rightarrow_g^* v \mid_{\mu'}$ for some v and μ' . Consequently,

$$\text{let } x^{\top} = t_2 \text{ in } t_1 \mid_{\mu} \rightarrow_g^* \text{let } x^{\top} = v \text{ in } t_1 \mid_{\mu'} \rightarrow_g t_1[x^{\top} \leftarrow v] \mid_{\mu'}$$

Therefore we have

$$\mathcal{E}_{\perp}(t) \mid_{\mathcal{E}(\mu)} \rightarrow^0 \mathcal{E}_{\perp}(\text{let } x^{\top} = v \text{ in } t_1) \mid_{\mathcal{E}(\mu)} \rightarrow \mathcal{E}_{\perp}(t_1)[x \leftarrow ()] \mid_{\mathcal{E}(\mu)}$$

By t -erasure definition, $\mathcal{E}_{\perp}(t_1)[x \leftarrow ()] = \mathcal{E}(t_2[x^{\top} \leftarrow v])$. Finally, by store preservation (Lemma 3), $\mu_{\perp} = \mu'_{\perp}$. Therefore, $\mathcal{E}(\mu) = \mathcal{E}(\mu')$, which allows us to conclude.

Case 2: t is $\text{let } x^{\perp} = t_2 \text{ in } t_1$. We have by hypothesis that t_1 is regular code. Then we also have that t_2 is regular too. If t_2 is a value, the result follows immediately. If t_2 is not a value, then the result follows from the results of GhostML soundness (safety and anti-monotony of statuses), using the induction hypothesis on t_2 , since t -erasure preserves the structure of regular terms.

Other cases are proved in a similar way. ■

Finally, we establish the second part of the main theorem:

Theorem 3 (Non-termination Preservation). If $\Sigma, \emptyset \vdash_g t : \tau, \perp$, (θ, ρ) holds and $t \mid_{\mu} \rightarrow_g \infty$, for some store μ such that $\Sigma \vdash_g \mu$, then $\mathcal{E}_{\perp}(t)$ also diverges, that is $\mathcal{E}_{\perp}(t) \mid_{\mathcal{E}(\mu)} \rightarrow \infty$.

Proof. By co-induction on $t \mid \mu \rightarrow_g \infty$. Observe that, since t diverges, t is not a value. Thus $\mathcal{E}_\perp(t)$ is not a value either. By well-typedness preservation (Theorem 1), $\mathcal{E}_\perp(t)$ is well-typed, with $\mathcal{E}(\Sigma), \emptyset \vdash \mathcal{E}_\perp(t) : \mathcal{E}_\perp(\tau), (\theta, \rho)$. We have $\mathcal{E}(\Sigma) \vdash \mathcal{E}(\mu)$, since $\Sigma \vdash_g \mu$. Therefore, by progress in MiniML, there exist some term q and some store ν such that $\mathcal{E}_\perp(t) \mid \mu \rightarrow q \mid \nu$. By backward simulation (Lemma 7), $t \mid \mu \rightarrow_g^{\geq 1} t' \mid \mu'$ with $q = \mathcal{E}_\perp(t')$ and $\nu = \mathcal{E}_\perp(\mu')$. Thus, the reduction $t \mid \mu \rightarrow_g \infty$ decomposes into

$$t \mid \mu \rightarrow_g^{\geq 1} t' \mid \mu' \rightarrow_g \infty$$

By preservation (Lemma 1), $\Sigma, \Gamma \vdash_g t' : \tau, \perp, (\theta, \rho)$ and $\Sigma \vdash_g \mu'$. Finally, by co-induction hypothesis, we get $\mathcal{E}_\perp(q) \mid \nu \rightarrow \infty$ and thus $\mathcal{E}_\perp(t) \mid \mathcal{E}(\mu) \rightarrow \infty$. ■

4 Implementation

Our method of ghost code handling 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 Hindley-Milner type polymorphism. Our approach to associate ghost status with variables and expressions, and not with types, makes this extension straightforward.

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, existing 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.

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.

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

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

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 implementations additionally stores the pure logical view of the queue as a list, in a 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)` 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 on-line 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 (aka 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 [9] and Reynolds [10], the notion of auxiliary variable was first introduced by Lucas in 1968 [11]. Since then, numerous authors have adapted this technique in various domains.

It is worth pointing out that some authors, in particular Kleymann [12] and Reynolds [10], 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.* [13] 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, the non-interference is ensured by

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


```

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.

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 [14] 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." [15]. In addition, a class can contain both ghost fields and regular fields. Dafny ensures termination of ghost code. However, apart from base types and operations, regular data types, such as arrays, cannot be reused in ghost code.

The property of non-interference of ghost code is not that different from non-interference is the context of information flow [16]. Indeed, one could see ghost code as high security level information and regular code as low level information, and non-interference precisely means that low level information is not computed out of high level one. Information flow can be checked using a type system [17] and proofs in that domain typically involve a bisimulation technique (though not necessarily through an erasure operation). Yet reducing the problem of ghost code analysis to information flow analysis still has to be investigated in details.

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, play 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. 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. `forall` loop construct in Dafny).

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. Plotkin, G.D.: Call-by-name, call-by-value and the lambda-calculus. Theor. Comput. Sci. **1**(2) (1975) 125–159
9. Jones, C.B., Roscoe, A., Wood, K.R.: Reflections on the Work of C.A.R. Hoare. 1st edn. Springer Publishing Company, Incorporated (2010)
10. Reynolds, J.C.: The craft of programming. Prentice Hall International series in computer science. Prentice Hall (1981)
11. Lucas, P.: Two constructive realizations of the block concept and their equivalence. Technical Report 25.085, IBM Laboratory, Vienna (June 1968)
12. Kleymann, T.: Hoare logic and auxiliary variables. Formal Asp. Comput. **11**(5) (1999) 541–566
13. 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
14. Schmaltz, S.: Towards the Pervasive Formal Verification of Multi-Core Operating Systems and Hypervisors Implemented in C. PhD thesis, Saarland University, Saarbrücken (2013)
15. Leino, K.R.M., Moskal, M.: Co-induction simply: Automatic co-inductive proofs in a program verifier. Technical Report MSR-TR-2013-49, Microsoft Research (July 2009) Available from <http://research.microsoft.com/pubs>.
16. Denning, D.E., Denning, P.J.: Certification of programs for secure information flow. Communications of the ACM **20**(2) (July 1977) 504–513
17. 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