

Global abstraction-safe marshalling via hash types

James J. Leifer Gilles Peskine Peter Sewell Keith Wansbrough

INRIA Rocquencourt

University of Cambridge

Problem

Consider inter-machine communication (or persistent storage):

(A)

```
...  
send (marshal (v : bool))
```

v →

(B)

```
...  
let y =  
  unmarshal (receive () : int list)
```

Problem

Consider inter-machine communication (or persistent storage):

(A)

```
...  
send (marshal (v : t ))
```

$v : t \rightarrow$

(B)

```
...  
let y =  
  unmarshal (receive () : t' )
```

A **dynamic type check** of $t = t'$ can ensure the safety of unmarshal.

Problem

Consider inter-machine communication (or persistent storage):

(A)

```
...  
send (marshal (v : t))
```

$\xrightarrow{v : t}$

(B)

```
...  
let y =  
  unmarshal (receive () : t')
```

A **dynamic type check** of $t = t'$ can ensure the safety of `unmarshal`.

But what if t and t' are ML-like abstract types, e.g.

```
t = UnbalancedBinaryTree.ty  
t' = BalancedBinaryTree.ty ?
```

Could just consider their concrete representation types to get type safety, but we want **abstraction safety** too.

Overview

- Examples: communication with abstract types
- Solution: hash types, compilation, and typing
- Theorems
- Conclusions and future work

An even counter: manifest signature

```
module EvenC = (  
  struct  
    type t      = int          (* the representation type *)  
    let start = 0  
    let up x   = x + 2  
    let get x  = x  
  end : EvenCSig)
```

```
EvenCSig =  
  sig  
    type t      = int          (* t is manifestly equal to int *)  
    val start   : t  
    val up      : t -> t  
    val get     : t -> int  
  end
```

An even counter: abstract signature

```
module EvenC = (  
  struct  
    type t      = int          (* the representation type *)  
    let start = 0  
    let up x   = x + 2  
    let get x = x  
  end : EvenCSig)
```

```
EvenCSig =  
  sig  
    type t          (* t is abstract *)  
    val start : t  
    val up    : t -> t  
    val get   : t -> int  
  end
```

Example: identical abstract types

(A)

```
module EvenC = (struct
  type t = int
  let start = 0
  let up x = x + 2
  let get x = x
end : EvenCSig)

let x = EvenC.start in
  send (marshal (x : EvenC.t))
```

(B)

```
module EvenC = (struct
  type t = int
  let start = 0
  let up x = x + 2
  let get x = x
end : EvenCSig)

let y =
  unmarshal (receive () : EvenC.t)
```

✓ succeed

Within a single program, two abstract types with the same definition would be different (ML generativity). Between programs, that's not what we want.

Example: concrete to abstract

(A)

...

```
let x = 3 in
  send (marshal (x : int))
```

(B)

```
module EvenC = (struct
  type t = int
  let start = 0
  let up x = x + 2
  let get x = x
end : EvenCSig)
```

```
let y =
  unmarshal (receive () : EvenC.t)
```

✗ fail

Allowing `unmarshal` to succeed would break (B)'s invariants.

Example: same external behaviour but different internal invariants

(A)

```
module EvenC = (struct
  type t = int
  let start = 0
  let up x = x + 1
  let get x = 2 * x
end : EvenCSig)

let x = EvenC.start in
  send (marshal (x : EvenC.t))
```

(B)

```
module EvenC = (struct
  type t = int
  let start = 0
  let up x = x + 2
  let get x = x
end : EvenCSig)

let y =
  unmarshal (receive () : EvenC.t)
```

✗ fail

Again, success would not respect (B)'s invariants.

Example: same internal invariants

(A) module EvenC = (struct type t = int let start = 0 let up x = 2 + x let get x = x end : EvenCSig) let x = EvenC.start in send (marshal (x : EvenC.t))	(B) module EvenC = (struct type t = int let start = 0 let up x = x + 2 let get x = x end : EvenCSig) let y = unmarshal (receive () : EvenC.t)
--	---

? maybe

Success would require a theorem prover to perform the verification (unrealistic) or a user-supplied coercion.

Summary of the main cases

Interface	Implementation	Desired behavior
same	same code	✓ succeed
same	same internal invariants	? maybe
same	same external behaviour but different internal invariants	✗ fail
same	different external behaviour	✗ fail
different	...	✗ fail
...	different representation types	✗ fail


How do we get the desired behaviour?

- For communication between **programs with identical sources**, it's easy to compare abstract types by their source-code names, e.g. `EvenC.t` would mean the same thing in all copies.
- However, for programs that share only some modules, that would be unsound.

How do we obtain **globally meaningful type names**?

Solution: we construct them from module ***hashes***.


(A)



`v : hash(struct ... end : sig ... end) .t`




(B)



Solution: hash types

- We can implement them with a cryptographic hash, e.g. md5 (compact fingerprint yet injective in practice).
- We freely look inside their structure in our typing rules, but never need to do this in the implementation.
- What exactly do we hash? A good candidate: abstract syntax trees of module definitions. But module dependencies require care.


(A)



$v : \text{hash}(\text{struct } \dots \text{end} : \text{sig } \dots \text{end}) . t$



(B)



1. Compile-time reduction: hash generation

```
module EvenC =  
  ( struct type t = int let start = 0 ... end  
    : sig type t val start : t ... end )  
send (marshal (EvenC.start : EvenC.t))
```

→_c

inlining EvenC

```
send (marshal (0 : h.t))
```

where $h = \text{hash}$ $\left(\begin{array}{l} \text{struct type } t = \text{int let start} = 0 \dots \text{end} \\ \text{: sig type } t \text{ val start : } t \dots \text{end} \end{array} \right)$

2. Compile-time reduction: module dependency (1/3)

```
module EvenC =  
  ( struct type t = int  
      let start = 0 ... end  
    : sig type t  
      val start : t ... end )  
  
module CleanC =  
  ( struct type s = EvenC.t * bool  
      let create = (EvenC.start, true) ... end  
    : sig type s  
      val create : s ... end )  
  
send (marshal (CleanC.create : CleanC.s))
```


2. Compile-time reduction: module dependency (2/3)

→_c

inlining `EvenC`

```
module CleanC =  
  ( struct type s = h.t * bool  
    let create = (0, true) ... end  
  : sig type s  
    val create : s ... end )  
send (marshal (CleanC.create : CleanC.s))
```

where

```
h = hash ( struct type t = int  
  let start = 0 ... end  
  : sig type t  
  val start : t ... end )
```

2. Compile-time reduction: module dependency (2/3)

→_c

inlining EvenC

```
module CleanC =  
  ( struct type s = h.t * bool  
    let create = (0, true) ... end  
  : sig type s  
    val create : s ... end )  
send (marshal (CleanC.create : CleanC.s))
```

where

```
h = hash ( struct type t = int  
  let start = 0 ... end  
  : sig type t  
  val start : t ... end )
```

2. Compile-time reduction: module dependency (3/3)

→_c

inlining CleanC

```
send (marshal ((0, true) : h'.s))
```

where

$$h = \text{hash} \left(\begin{array}{l} \text{struct} \quad \text{type } t = \text{int} \\ \quad \quad \text{let } \text{start} = 0 \quad \dots \quad \text{end} \\ : \text{sig} \quad \text{type } t \\ \quad \quad \text{val } \text{start} : t \quad \dots \quad \text{end} \end{array} \right)$$
$$h' = \text{hash} \left(\begin{array}{l} \text{struct} \quad \text{type } s = h.t * \text{bool} \\ \quad \quad \text{let } \text{create} = (0, \text{true}) \quad \dots \quad \text{end} \\ : \text{sig} \quad \text{type } s \\ \quad \quad \text{val } \text{create} : s \quad \dots \quad \text{end} \end{array} \right)$$

3. Compile-time reduction: coloured brackets

```
module EvenC =  
  ( struct type t = int let start = 0 ... end  
    : sig type t      val start : t ... end )  
send (marshal (EvenC.start : EvenC.t))
```

→_c

inlining EvenC

```
send (marshal ( 0 : h.t ))
```

where

```
h = hash ( struct type t = int let start = 0 ... end  
           : sig type t      val start : t ... end )
```

Coloured brackets are adapted from [Zdancewic, Grossman, & Morrisett]

3. Compile-time reduction: coloured brackets

```
module EvenC =  
  ( struct type t = int let start = 0 ... end  
    : sig type t      val start : t ... end )  
send (marshal (EvenC.start : EvenC.t))
```

→_c

inlining EvenC

```
send (marshal ([0] h.t : h.t))
```

where

```
h = hash ( struct type t = int let start = 0 ... end  
           : sig type t      val start : t ... end )
```

Coloured brackets are adapted from [Zdancewic, Grossman, & Morrisett]

The calculus

- call-by-value lambda-calculus;
- second-class, first-order modules;
- communication and parallel composition;
- marshal and unmarshal;
- hashes in the type grammar:

$$\begin{array}{l} \mathbb{T} ::= \dots \\ \quad | \quad h.t \quad \text{(not in user source code)} \end{array}$$

- coloured brackets in the expression grammar:

$$\begin{array}{l} e ::= \dots \\ \quad | \quad [e]_h^T \quad \text{(not in user source code)} \end{array}$$

Type equality ($E \vdash_h T_0 == T_1$)

- **singleton kind** equations for module typing [Harper & Lillibridge];
- plus **hash transparency** when inside coloured brackets:

$$\frac{E \vdash_h \text{ok}}{E \vdash_h h.t == T} \quad \text{if } h = \text{hash} \left(\begin{array}{l} \text{struct type } t = T \quad \dots \quad \text{end} \\ : \text{ sig type } t \quad \dots \quad \text{end} \end{array} \right)$$

Coloured brackets

- determine where hash transparency occurs:

$$\frac{E \vdash_h e : T}{E \vdash_{h'} [e]_h^T : T}$$

Theorems

- **Type preservation, progress:** for compile-time and run-time reduction. Thanks to brackets, this includes (informally) **abstraction preservation**.
- **Type coincidence:** ML type equivalence coincides with unmarshal-time syntactic comparison of hash types.
- **Erasure:** after compilation, erasure of all coloured brackets (except in hashes) yields identical run-time behaviour.

Subtleties: handling dependent signatures and tracking colours.
(We optimise proofs with a rigorous meta-notion of “similar case”.)

Conclusions and future work

Hashing modules provides a meaningful way of comparing abstract types that are defined in independently compiled distinct programs: as a result, the behaviour we sought “just works”.

What's next?

- **ML**: multiple type and term fields, polymorphism, functors, nested modules;
- **Beyond**: subtyping, coercions and versioning, dynamic binding for local resources;
- **Implementation**: Jocaml and Ocaml, applications to safe name servers, channels, persistent stores.

Theorems in detail...

Theorem: erasure

After compilation, erasure of all coloured brackets (except in hashes) yields identical run-time behaviour.

Let `erase` be the erasure function.

Let $\xrightarrow{\text{uncol}}$ be the corresponding run-time reduction relation.

- If $\text{nil} \vdash_{\text{ho}} e:T$ and $e \xrightarrow{\text{ho}} e'$ then $\text{erase}(e) \xrightarrow{\text{uncol}}^{\leq 1} \text{erase}(e')$.
- If $\text{nil} \vdash_{\text{ho}} e:T$ and $\text{erase}(e) \xrightarrow{\text{uncol}} e_0$
then there exists e' such that $\text{erase}(e') = e_0$ and $e \xrightarrow{\text{ho}}^{\geq 1} e'$.

As we said, $\xrightarrow{\text{uncol}}$ does not preserve typing.

Theorem: type coincidence

ML type equivalence coincides with unmarshal-time syntactic comparison of hash types.

Consider the following sequence of module definitions:

$$D._ = \text{module } U_1 = M_1:S_1 \text{ in } \dots \text{module } U_n = M_n:S_n \text{ in } _$$

Let σ_D be the substitution induced by compilation of the modules.

Suppose that no two modules have the same hash.

Then:

$$\begin{array}{l} U_1:S_1, \dots, U_n:S_n \vdash_{\bullet} T_0 == T_1 \quad \iff \quad \sigma_D T_0 = \sigma_D T_1 \\ \text{static typing} \quad \iff \quad \text{check performed by unmarshal} \end{array}$$