

Imperative HOL – a lightweight framework for imperative data structures in Isabelle/HOL

March 13, 2025

Imperative HOL is a lightweight framework for reasoning about imperative data structures in *Isabelle/HOL* [2]. Its basic ideas are described in [1]. However their concrete realisation has changed since, due to both extensions and refinements. Therefore this overview wants to present the framework “as it is” by now. It focusses on the user-view, less on matters of construction. For details study of the theory sources is encouraged.

1 A polymorphic heap inside a monad

Heaps (*heap*) can be populated by values of class *heap*; HOL’s default types are already instantiated to class *heap*. Class *heap* is a subclass of *countable*; see theory *Countable* for ways to instantiate types as *countable*.

The heap is wrapped up in a monad *'a Heap* by means of the following specification:

datatype *'a Heap* = *Heap.Heap* (*heap* \Rightarrow (*'a* \times *heap*) *option*)

Unwrapping of this monad type happens through

execute
execute (*Heap.Heap* *f*) = *f*

This allows for equational reasoning about monadic expressions; the fact collection *execute_simps* contains appropriate rewrites for all fundamental operations.

Primitive fine-granular control over heaps is available through rule *Heap-cases*:

$$(\bigwedge x\ h'.\ \text{execute } f\ h = \text{Some } (x, h') \implies P) \implies (\text{execute } f\ h = \text{None} \implies P) \implies P$$

Monadic expression involve the usual combinators:

return
 $(\gg=)$
raise

This is also associated with nice monad do-syntax. The *string* argument to *raise* is just a codified comment.

Among a couple of generic combinators the following is helpful for establishing invariants:

assert
assert $P\ x = (\text{if } P\ x \text{ then return } x \text{ else raise STR "assert"})$

2 Relational reasoning about *Heap* expressions

To establish correctness of imperative programs, predicate

effect

provides a simple relational calculus. Primitive rules are *effectI* and *effectE*, rules appropriate for reasoning about imperative operations are available in the *effect-intros* and *effect-elim*s fact collections.

Often non-failure of imperative computations does not depend on the heap at all; reasoning then can be easier using predicate

success

Introduction rules for *success* are available in the *success-intro* fact collection.

execute, *effect*, *success* and $(\gg=)$ are related by rules *execute-bind-success*, *success-bind-executeI*, *success-bind-effectI*, *effect-bindI*, *effect-bindE* and *execute-bind-eq-SomeI*.

3 Monadic data structures

The operations for monadic data structures (arrays and references) come in two flavours:

- Operations on the bare heap; their number is kept minimal to facilitate proving.
- Operations on the heap wrapped up in a monad; these are designed for executing.

Provided proof rules are such that they reduce monad operations to operations on bare heaps.

Note that HOL equality coincides with reference equality and may be used as primitive executable operation.

3.1 Arrays

Heap operations:

Array.alloc
Array.present
Array.get
Array.set
Array.length
Array.update
 (==)

Monad operations:

Array.new
Array.of-list
Array.make
Array.len
Array.nth
Array.upd
Array.map-entry
Array.swap
Array.freeze

3.2 References

Heap operations:

Ref.alloc
Ref.present
Ref.get
Ref.set
 (==)

Monad operations:

ref
Ref.lookup
Ref.update
Ref.change

4 Code generation

Imperative HOL sets up the code generator in a way that imperative operations are mapped to suitable counterparts in the target language. For *Haskell*, a suitable *ST* monad is used; for *SML*, *Ocaml* and *Scala* unit values ensure that the evaluation order is the same as you would expect from the original monadic expressions. These units may look cumbersome; the target language variants *SML-imp*, *Ocaml-imp* and *Scala-imp* make some effort to optimize some of them away.

5 Some hints for using the framework

Of course a framework itself does not by itself indicate how to make best use of it. Here some hints drawn from prior experiences with Imperative HOL:

- Proofs on bare heaps should be strictly separated from those for monadic expressions. The first capture the essence, while the latter just describe a certain wrapping-up.
- A good methodology is to gradually improve an imperative program from a functional one. In the extreme case this means that an original functional program is decomposed into suitable operations with exactly one corresponding imperative operation. Having shown suitable correspondence lemmas between those, the correctness prove of the whole imperative program simply consists of composing those.
- Whether one should prefer equational reasoning (fact collection *execute-simps* or relational reasoning (fact collections *effect-intros* and *effect-elim*s) depends on the problems to solve. For complex expressions or expressions involving binders, the relation style is usually superior but requires more proof text.
- Note that you can extend the fact collections of Imperative HOL yourself whenever appropriate.

References

- [1] L. Bulwahn, A. Krauss, F. Haftmann, L. Erkök, and J. Matthews. Imperative functional programming with Isabelle/HOL. In O. A. Mohamed, C. Muñoz, and S. Tahar, editors, *TPHOLs '08: Proceedings of the 21th International Conference on Theorem Proving in Higher Order Logics*, volume 5170 of *Lecture Notes in Computer Science*, pages 352–367. Springer-Verlag, 2008.

- [2] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer-Verlag, 2002.