

Effect polymorphism in higher-order logic (Proof Pearl)

Andreas Lochbihler

Institute of Information Security, Department of Computer Science, ETH Zurich, Switzerland

Abstract. The notion of a *monad* cannot be expressed within higher-order logic (HOL) due to type system restrictions. We show that if a monad is used with values of only one type, this notion *can* be formalised in HOL. Based on this idea, we develop a library of effect specifications and implementations of monads and monad transformers. Hence, we can abstract over the concrete monad in HOL definitions and thus use the same definition for different (combinations of) effects. We illustrate the usefulness of effect polymorphism with a monadic interpreter.

1 Introduction

Monads have become a standard way to write effectful programs in pure functional languages [25]. In proof assistants, they provide a widely-used abstraction for modelling and reasoning about effects [3,4,14,17]. Abstractly, a monad consists of a type constructor τ and two polymorphic operations, `return` $:: \alpha \Rightarrow \alpha \tau$ for embedding values and `bind` $:: \alpha \tau \Rightarrow (\alpha \Rightarrow \beta \tau) \Rightarrow \beta \tau$ for sequencing (written $\gg=$ infix), satisfying three monad laws:

1. $(m \gg= f) \gg= g = m \gg= (\lambda x. f x \gg= g)$
2. $\text{return } x \gg= f = f x$
3. $m \gg= \text{return} = m$

Yet, the notion of a monad cannot be expressed as a formula in higher-order logic (HOL) [8] as there are no type constructor variables like τ in HOL and the sequencing operation `bind` occurs with three different type instances in the first law. Thus, only concrete monad instances have been used to model side effects of HOL functions. In fact, monad definitions for different effects abound in HOL, e.g., a state-error monad [3], non-determinism with errors and divergence [14], probabilistic choice [4], and probabilistic resumptions with errors [17]. Each of these formalisations fixes τ to a particular type (constructor) and develops its own reasoning infrastructure. This approach achieves *value polymorphism*, i.e., one monad can be used with varying types of values, but not *effect polymorphism* where one function can be used with different monads.

In this paper, we give up value polymorphism in favour of effect polymorphism. The idea is to fix the type of values to some type α_0 . Then, the monad type constructor τ is applied only to α_0 , which an ordinary HOL type variable μ can represent. So, the monad operations have the HOL types `return` $:: \alpha_0 \Rightarrow \mu$ and `bind` $:: \mu \Rightarrow (\alpha_0 \Rightarrow \mu) \Rightarrow \mu$. This notion of a monad can be formalised within HOL. In detail, we present an Isabelle/HOL library (available online [18]) for different monadic effects and their algebraic specification. All effects are also implemented as value-monomorphic monads and monad transformers. Using Isabelle's module system [1], function definitions can be made abstractly and later specialised to several concrete monads. As our running example, we formalise and reason about a monadic interpreter for a small language. The library has been used in a larger project to define and reason about parsers and serialisers for security protocols.

Contributions. We show the advantages of trading in value polymorphism for effect polymorphism. First, HOL functions with effects can be defined in an abstract monadic setting (§2) and reasoned about in the style of Gibbons and Hinze [6]. This preserves the level of abstraction that the monad notion provides. As the definitions need not commit to a concrete monad, we can use them in richer effect contexts, too—simply by combining our modular effect specifications. When a concrete monad instance is needed, it can be easily obtained by interpretation using Isabelle’s module system.

Second, as HOL can express the notion of a value-monomorphic monad, we have also formalised several monad transformers [15,21] in HOL (§3). Thus, there is no need to define the monad and derive the reasoning principles for each combination of effects, as is current practice with value polymorphism. Instead, it suffices to formalise every effect only once as a transformer and combine them modularly.

Third, relations between different instances can be proven using the theory of representation independence (§4) as supported by the Transfer package [10]. This makes it possible to switch in the middle of a bigger proof from a complicated monad to a simpler one.

2 Abstract Value-Monomorphic Monads in HOL

In this section, we formalise value-monomorphic monads and monad transformers for several types of effects. A monadic interpreter for an arithmetic language will be used throughout as a running example. The language, adapted from Nipkow and Klein [22], consists of integer constants, variables, addition, and division.

```
datatype v exp = Const int | Var v | (v exp) ⊕ (v exp) | (v exp) ⊗ (v exp)
```

We formalise the concept of a monad using Isabelle’s module system of locales [1]. The locale monad below fixes the two monad operations return and bind (written infix as $\gg\equiv$) and assumes that the monad laws hold. It will collect definitions of functions, which use the monad operations, and theorems about them, whose proofs can use the monad laws. Every locale also defines a predicate of the same name that collects all the assumptions. When a user interprets the locale with more concrete operations and has discharged the assumptions for these operations, every definition and theorem inside the locale context is specialised to these operations. Although the type of values is a type variable α , α is fixed inside the locale. Instantiations may still replace α with any other HOL type. In other words, the locale monad formalises a *monomorphic* monad, but leaves the type of values unspecified. As usual, $m \gg m'$ abbreviates $m \gg\equiv (\lambda_. m')$.

```
locale monad = fixes return ::  $\alpha \Rightarrow \mu$  and bind ::  $\mu \Rightarrow (\alpha \Rightarrow \mu) \Rightarrow \mu$  (infixr  $\gg\equiv$ )
  assumes BIND-ASSOC:  $(m \gg\equiv f) \gg\equiv g = m \gg\equiv (\lambda x. f x \gg\equiv g)$ 
  and RETURN-BIND:  $\text{return } x \gg\equiv f = f x$ 
  and BIND-RETURN:  $x \gg\equiv \text{return} = x$ 
```

Monads become useful only when effect-specific operations are available. In the remainder of this section, we formalise monadic operations for different types of effects and their properties. For each effect, we introduce a new locale in Isabelle that extends the locale monad, fixes the new operations, and specifies their properties. A locale extension inherits parameters and assumptions. This leads to a modular design: if several effects are needed, one merely combines the relevant locales in a multi-extension.

2.1 Failure and Exception

Failures are one of the simplest effects and widely used. A failure aborts the computation immediately. The locale `monad-fail` given below formalises the failure effect `fail :: μ`. It assumes that a failure propagates from the left hand side of `bind`. In contrast, there is no assumption about how `fail` behaves on the right hand side. Otherwise, if `monad-fail` also assumed $m \gg= (\lambda _. \text{fail}) = \text{fail}$, then `fail` would undo any effect of m . Although the standard implementation of failures using the option type satisfies this additional law, many other monad implementations do not, e.g., resumptions. Note that there is no need to delay the evaluation of `fail` in HOL because HOL has no execution semantics.

```
locale monad-fail = monad + fixes fail :: μ
  assumes FAIL-BIND: fail >>= f = fail
```

As a first example, we define the monadic interpreter `eval :: (ν ⇒ μ) ⇒ ν exp ⇒ μ` for arithmetic expressions by primitive recursion using these abstract monad operations inside the locale `monad-fail`.¹ The first argument is an interpretation function $E :: ν ⇒ μ$ for the variables. The evaluation fails when a division by zero occurs.

```
primrec (in monad-fail) eval :: (ν ⇒ μ) ⇒ ν exp ⇒ μ where
  eval E (Const i) = return i
| eval E (Var x)   = E x
| eval E (e1 ⊕ e2) = eval E e1 >>= (λ i1. eval E e2 >>= (λ i2. return (i1 + i2)))
| eval E (e1 ⊙ e2) =
  eval E e1 >>= (λ i1. eval E e2 >>= (λ i2. if i2 = 0 then fail else return (i1 div i2)))
```

Note that evaluating a variable can have an effect $μ$, which is necessary to obtain a compositional interpreter. Let `subst :: (ν ⇒ ν' exp) ⇒ ν exp ⇒ ν' exp` be the substitution function for `exp`. That is, `subst σ e` replaces every `Var x` in e with $σ x$. Then, the following compositionality statement holds (proven by induction on e and term rewriting with the definitions), where function composition $∘$ is defined as $(f ∘ g)(x) = f (g x)$.

```
lemma COMPOSITIONALITY: eval E (subst σ e) = eval (eval E ∘ σ) e
  by induction simp-all
```

We refer to failures as exceptions whenever there is an operator `catch :: μ ⇒ μ ⇒ μ` to handle them. Following Gibbons and Hinze [6], the locale `monad-catch` assumes that `catch` and `fail` form a monoid and that returns are not handled. It inherits `FAIL-BIND` and the monad laws by extending the locale `monad-fail`. No properties about `catch` and `bind` are assumed because in general exception handling does not distribute over sequencing.

```
locale monad-catch = monad-fail + fixes catch :: μ ⇒ μ ⇒ μ
  assumes FAIL-CATCH: catch fail m = m
  and CATCH-FAIL:    catch m fail = m
  and CATCH-CATCH:   catch (catch m1 m2) m3 = catch m1 (catch m2 m3)
  and RETURN-CATCH:  catch (return x) m = return x
```

¹ Type variables that appear in the signature of locale parameters are fixed for the whole locale. In particular, the value type $α$ cannot be instantiated inside the locale `monad` or its extension `monad-fail`. The interpreter `eval`, however, returns `ints`. For this reason, `eval` is defined in an extension of `monad-fail` that merely specialises $α$ to `int`. For readability, we usually omit this detail in this paper.

2.2 State

Stateful computations use operations to read (`get`) and replace (`put`) the state of type σ . In a value-polymorphic setting, $\text{get} :: \sigma \rightarrow \tau$ and $\text{put} :: \sigma \Rightarrow \text{unit } \tau$ are usually computations that return the state or $()$ inhabiting the singleton type `unit`. Without value-polymorphism, these types cannot be formalised in the HOL setting because we cannot apply τ to different value types. Instead, our operations additionally take a continuation: $\text{get} :: (\sigma \Rightarrow \mu) \Rightarrow \mu$ and $\text{put} :: \sigma \Rightarrow \mu \Rightarrow \mu$. In a value-polymorphic setting, both signatures are equivalent. Passing the continuation `return` as in `get return` and $\lambda s. \text{put } s (\text{return } ())$ yields the conventional operations. Conversely, our operations `get f` and `put s m` can be implemented as $\text{get} \ggg f$ and $\text{put } s \ggg m$ using conventional `get` and `put`. The locale `monad-state` collects the properties `get` and `put` must satisfy:

```

locale monad-state = monad + fixes get :: ( $\sigma \Rightarrow \mu$ )  $\Rightarrow \mu$  and put ::  $\sigma \Rightarrow \mu \Rightarrow \mu$ 
  assumes PUT-GET: put s (get f) = put s (f s)
    and GET-GET:  get ( $\lambda s. \text{get } (f s)$ ) = get ( $\lambda s. f s s$ )
    and PUT-PUT:  put s (put s' m) = put s' m
    and GET-PUT:  get ( $\lambda s. \text{put } s m$ ) = m
    and GET-CONST: get ( $\lambda \_ . m$ ) = m
    and BIND-GET: get f  $\ggg$  g = get ( $\lambda s. f s \ggg g$ )
    and BIND-PUT: put s m  $\ggg$  f = put s (m  $\ggg$  f)

```

The first four assumptions adapt Gibbons' and Hinze's axioms for the state operations [6] to the new signature. The fifth, `GET-CONST`, additionally specifies that `get` can be discarded if the state is not used. The last two assumptions, `BIND-GET` and `BIND-PUT`, demand that `get` and `put` distribute over `bind`. In the conventional value-polymorphic setting, where the continuations are applied using `bind`, these two are subsumed by the monad laws. In the remainder of this paper, `get` and `put` always take continuations.

A state update function `update` can be implemented abstractly for all state monads. Like `put`, `update` takes a continuation m .

```

definition (in monad-state) update :: ( $\sigma \Rightarrow \sigma$ )  $\Rightarrow \mu \Rightarrow \mu$  where
  update f m = get ( $\lambda s. \text{put } (f s) m$ )

```

The expected properties of `update` can be derived from `monad-state`'s assumptions by term rewriting. For example,

```

lemma UPDATE-ID: update id m = m
  by (simp add: UPDATE-DEF GET-PUT)

lemma UPDATE-UPDATE: update f (update g m) = update (g  $\circ$  f) m
  by (simp add: UPDATE-DEF PUT-GET PUT-PUT)

lemma UPDATE-BIND: update f m  $\ggg$  g = update f (m  $\ggg$  g)
  by (simp add: UPDATE-DEF BIND-GET BIND-PUT)

```

As an example, we implement a memoisation operator `memo` using the state operations. To that end, the state must be refined to a lookup table, which we model as a map of type $\beta \rightarrow \alpha = \beta \Rightarrow \alpha$ option. The definition uses the function $\lambda t. t(x \mapsto y)$ that takes a map t and updates it to associate x with y , leaving the other associations as they are; formally, $t(x \mapsto y) = (\lambda x'. \text{if } x = x' \text{ then Some } y \text{ else } t x')$.

definition (in monad-state) memo :: ($\beta \Rightarrow \mu$) \Rightarrow $\beta \Rightarrow \mu$ where
 memo $f x =$ get ($\lambda table.$
 case $table x$ of Some $y \Rightarrow$ return y
 | None $\Rightarrow f x \gg\equiv (\lambda y. update (\lambda t. t(x \mapsto y)) (return y))$)

A memoisation operator should satisfy three important properties. First, it should evaluate the memoised function at most on the given argument, not on others. This can be expressed as a congruence rule, which holds independently of the monad laws by definition:

lemma MEMO-CONG: $f x = g x \longrightarrow memo f x = memo g x$

Second, memoisation should be idempotent, i.e., if a function is already being memoised, then there is no point in memoising it once more.

lemma MEMO-IDEM: $memo (memo f) x = memo f x$

The mechanised proof of MEMO-IDEM in Isabelle needs only two steps, which are justified by term rewriting with the properties of the monad operations and the case operator. Every assumption about get and put except GET-PUT is needed.

Third, the memoisation operator should indeed evaluate f on x at most once. As memo $f x$ memoises only the result of $f x$, but not the effect of evaluating $f x$, the next lemma captures this correctness property. Its proof is similar to MEMO-IDEM's.

lemma CORRECT: $memo f x \gg\equiv (\lambda a. memo f x \gg\equiv g a) = memo f x \gg\equiv (\lambda a. g a a)$

2.3 Probabilistic Choice

Randomised computations are built from an operation \clubsuit for probabilistic choice. The probabilities are specified using probability mass functions (type π pmf) [7], i.e., discrete probability distributions. Binary probabilistic choice, which is often used in the literature [5,6,24], is less general as it leads to finite distributions. Continuous distributions would work, too, but they would clutter the theorems and proofs with measurability conditions.

Like the state operations, $\clubsuit :: \pi \text{ pmf} \Rightarrow (\pi \Rightarrow \mu) \Rightarrow \mu$ takes a continuation to separate the type of probabilistic choices π from the type of values. The locale monad-prob assumes the following properties, where $\text{supp } p$ denotes the support of p :

- sampling from the one-point distribution $\text{dirac } x$ has no effect (SAMPLE-DIRAC),
- sequencing bind_{pmf} in the probability monad yields sequencing (SAMPLE-BIND),
- sampling can be discarded if the result is unused (SAMPLE-CONST),
- sampling from independent distributions commutes (SAMPLE-COMM, independence is formalised by p and q not taking y and x as an argument, respectively.)
- sampling calls the continuation only on values in p 's support (SAMPLE-CONG), and
- sampling distributes over both sides of bind (BIND-SAMPLE₁, BIND-SAMPLE₂).

locale monad-prob = monad + fixes $\clubsuit :: \pi \text{ pmf} \Rightarrow (\pi \Rightarrow \mu) \Rightarrow \mu$

assumes SAMPLE-DIRAC: $\clubsuit (\text{dirac } x) f = f x$
 and SAMPLE-BIND: $\clubsuit (\text{bind}_{\text{pmf}} p f) g = \clubsuit p (\lambda x. \clubsuit (f x) g)$
 and SAMPLE-CONST: $\clubsuit p (\lambda _. m) = m$
 and SAMPLE-COMM: $\clubsuit p (\lambda x. \clubsuit q (f x)) = \clubsuit q (\lambda y. \clubsuit p (\lambda x. f x y))$
 and SAMPLE-CONG: $(\forall x \in \text{supp } p. f x = g x) \longrightarrow \clubsuit p f = \clubsuit p g$
 and BIND-SAMPLE₁: $\clubsuit p f \gg\equiv g = \clubsuit p (\lambda x. f x \gg\equiv g)$
 and BIND-SAMPLE₂: $m \gg\equiv (\lambda x. \clubsuit p (f x)) = \clubsuit p (\lambda y. m \gg\equiv (\lambda x. f x y))$

2.4 Combining Abstract Monads

Formalising monads in this abstract way has the advantage that the different effects can be easily combined. In the running example, suppose that the variables represent independent random variables. Then, expressions are probabilistic computations and evaluation computes the joint probability distribution. For example, if x_1 and x_2 represent coin flips with 1 representing heads and 0 tails, then $\text{Var } x_1 \oplus \text{Var } x_2$ represents the probability distribution of the number of heads.

Here is a first attempt. Let $X :: v \Rightarrow \text{int pmf}$ specify the distribution $X x$ for each random variable x . Combining the locales for failures and probabilistic choices, we let the variable environment do the sampling, where $\text{sample-var } X x = \mathfrak{c} (X x)$ return:

```
locale monad-fail-prob = monad-fail + monad-prob
```

```
definition (in monad-fail-prob) wrong :: (v  $\Rightarrow$  int pmf)  $\Rightarrow$  v exp  $\Rightarrow$   $\mu$  where
  wrong X e = eval (sample-var X) e
```

As the name suggests, `wrong` does not achieve what we intended. If a variable occurs multiple times in e , say $e = \text{Var } x \oplus \text{Var } x$, then `wrong X e` samples x afresh for each occurrence. So, if $X x = \text{uniform } \{0, 1\}$, i.e., x is a coin flip, `wrong X e` computes the probability distribution given by $0 \mapsto 1/4, 1 \mapsto 1/2, 2 \mapsto 1/4$ instead of $0 \mapsto 1/2, 2 \mapsto 1/2$. Clearly, we should sample every variable at most once. Memoising the variable evaluation achieves that. So, we additionally need state operations.

```
locale monad-fail-prob-state = monad-fail-prob + monad-state +
  assumes SAMPLE-GET:  $\mathfrak{c} p (\lambda x. \text{get } (f x)) = \text{get } (\lambda s. \mathfrak{c} p (\lambda x. f x s))$ 
```

```
definition (in monad-fail-prob-state) lazy :: (v  $\Rightarrow$  int pmf)  $\Rightarrow$  v exp  $\Rightarrow$   $\mu$  where
  lazy X e = eval (memo (sample-var X)) e
```

The interpreter `lazy` samples a variable only when needed. For example, in $e_0 = (\text{Const } 1 \odot \text{Const } 0) \oplus \text{Var } x_0$, the division by zero makes the evaluation fail before x_0 is sampled.

The locale `monad-fail-prob-state` adds an assumption that \mathfrak{c} distributes over `get`. Such distributivity assumptions are typically needed because of the continuation parameters, which break the separation between effects and sequencing. Their format is as follows: If two operations f_1 and f_2 with continuations do not interact, then we assume $f_1 (\lambda x. f_2 (g x)) = f_2 (\lambda y. f_1 (\lambda x. g x y))$. Sometimes, such assumptions follow from existing assumptions. For example, `SAMPLE-PUT` follows from `BIND-SAMPLE2` and $\text{put } s m = \text{put } s (\text{return } x) \gg m$ for all x . A similar law holds for `update`.

```
lemma SAMPLE-PUT:  $\mathfrak{c} p (\lambda x. \text{put } s (f x)) = \text{put } s (\mathfrak{c} p f)$ 
```

In contrast, `SAMPLE-GET` does not follow from the other assumptions due to the restriction to monomorphic values. The state of type σ , which `get` passes to its continuation, may carry more information than a value can hold. Indeed, in the case of `lazy`, the type `int` of values is countable, but the state type $v \rightarrow \text{int}$ is not if the type of variables is infinite. As `put` passes no information to its continuation, `put`'s continuation can be pushed into `bind` as shown above. Still, `put` needs its continuation; otherwise, it would have to create a return value out of nothing, which would cause problems later (§4). Moreover, there is no need to explicitly specify how `fail` interacts with `get` and \mathfrak{c} as $\text{get } (\lambda _. \text{fail}) = \text{fail}$ and $\mathfrak{c} p (\lambda _. \text{fail}) = \text{fail}$ are special cases of `GET-CONST` and `SAMPLE-CONST`.

Instead of lazy sampling, we can also sample all variables eagerly. Let $\text{vars } e$ return the (finite) set of variables in e . Then, the interpreter eager with eager sampling is defined as follows (all three definitions live in the locale `monad-fail-prob-state`):

definition `sample-vars` :: $(v \Rightarrow \text{int pmf}) \Rightarrow v \text{ set} \Rightarrow \mu \Rightarrow \mu$ where
`sample-vars` $X A m = \text{fold } (\lambda x m. \text{memo } (\text{sample-var } X) x \gg m) m A$

definition `lookup` :: $v \Rightarrow \mu$ where
`lookup` $x = \text{get } (\lambda s. \text{case } s \ x \ \text{of } \text{None} \Rightarrow \text{fail} \mid \text{Some } i \Rightarrow \text{return } i)$

definition `eager` :: $(v \Rightarrow \text{int pmf}) \Rightarrow v \text{ exp} \Rightarrow \mu$ where
`eager` $X e = \text{sample-vars } X (\text{vars } e) (\text{eval lookup } e)$

where `fold` is the fold operator for finite sets [23]. The operator `fold` f requires that the folding function f is left-commutative, i.e., $f \ x \ (f \ y \ z) = f \ y \ (f \ x \ z)$ for all x , y , and z . In our case, $f = \lambda x m. \text{memo } (\text{sample-var } X) x \gg m$ is left-commutative by the following lemma about `memo` whose assumptions `sample-var` X satisfies by `RETURN-BIND`, `BIND-SAMPLE1`, `BIND-SAMPLE2`, and `SAMPLE-GET`. Moreover, by `CORRECT`, it is also idempotent, i.e., $f \ x \circ f \ x = f \ x$.

lemma `MEMO-COMMUTE` :

$$\begin{aligned} & (\forall m \ x \ g. m \gg (\lambda a. f \ x \gg g \ a) = f \ x \gg (\lambda b. m \gg (\lambda a. g \ a \ b))) \\ \longrightarrow & (\forall x \ g. \text{get } (\lambda s. f \ x \gg g \ s) = f \ x \gg (\lambda a. \text{get } (\lambda s. g \ s \ a))) \\ \longrightarrow & \text{memo } f \ x \gg (\lambda a. \text{memo } f \ y \gg (\lambda b. g \ a \ b)) = \\ & \text{memo } f \ y \gg (\lambda b. \text{memo } f \ x \gg (\lambda a. g \ a \ b)) \end{aligned}$$

This lemma and `CORRECT` illustrate the typical form of monadic statements. The assumptions and conclusions take a continuation g for the remainder of the program. This way, the statements are easier to apply because they are in normal form with respect to `BIND-ASSOC`. This observation also holds in a value-polymorphic setting.

Now, the question is whether eager and lazy sampling are equivalent. In general, the answer is no. For example, for e_0 from above, `eager` $X \ e_0$ samples and memoises the variable x_0 , but `lazy` $X \ e_0$ does not. Thus, there are contexts that distinguish the two. If we extend `monad-fail-prob-state` with exception handling from `monad-catch` such that

`CATCH-GET`: $\text{catch } (\text{get } f) \ m_2 = \text{get } (\lambda s. \text{catch } (f \ s) \ m_2)$
`CATCH-PUT`: $\text{catch } (\text{put } s \ m) \ m_2 = \text{put } s \ (\text{catch } m \ m_2)$

then the two can be distinguished:

$$\begin{aligned} & \text{catch } (\text{lazy } X \ e_0) \ (\text{lookup } x_0) = \text{fail} \\ & \text{catch } (\text{eager } X \ e_0) \ (\text{lookup } x_0) = \text{memo } (\text{sample-var } X) \ x_0 \end{aligned}$$

In contrast, if we assume that failures erase state updates, then the two *are* equivalent:

theorem `LAZY-EAGER` : $(\forall s. \text{put } s \ \text{fail} = \text{fail}) \longrightarrow \text{lazy } X \ e = \text{eager } X \ e$

Proof. The proof consists of three steps proven by induction on e . First, by idempotence and left-commutativity, `sample-vars` $X \ V$ commutes with `lazy` $X \ e$ for any finite V :

$$\forall g. \text{sample-vars } X \ V (\text{lazy } X \ e \gg g) = \text{lazy } X \ e \gg (\lambda i. \text{sample-vars } X \ V (g \ i)) \quad (1)$$

Here, `put` $s \ \text{fail} = \text{fail}$ ensures that all state updates are lost if a division by zero occurs. The next two steps will use (1) in the inductive cases for \oplus and \otimes to bring together the sampling of the variables and the evaluation of the subexpressions. Second,

$$\text{lazy } X e \ggg g = \text{sample-vars } X (\text{vars } e) (\text{lazy } X e \ggg g) \quad (2)$$

shows that the sampling can be done first, which holds by CORRECT. Finally,

$$\text{sample-vars } X V (\text{lazy } X e \ggg g) = \text{sample-vars } X V (\text{eval lookup } e \ggg g) \quad (3)$$

holds for any finite set V with $\text{vars } e \subseteq V$. Here, $\text{Var } x$ is the interesting case, which follows from $\forall g. \text{memo } f x \ggg (\lambda i. \text{lookup } x \ggg g i) = \text{memo } f x \ggg (\lambda i. g i i)$ and CORRECT. Taking $V = \text{vars } e$ and $g = \text{return}$, (2) and (3) prove the lemma. \square

In §3.5, we show that some monads satisfy LAZY-EAGER's assumption, but not all.

2.5 Further Abstract Monads

Apart from exceptions, state, and probabilistic choice, we have formalised effect specifications for non-deterministic choice $\text{alt} :: \mu \Rightarrow \mu \Rightarrow \mu$, the reader and writer monads with $\text{ask} :: (\rho \Rightarrow \mu) \Rightarrow \mu$ and $\text{tell} :: \omega \Rightarrow \mu \Rightarrow \mu$, and resumptions with $\text{pause} :: \circ \Rightarrow (t \Rightarrow \mu) \Rightarrow \mu$. We do not present them in detail as the examples in this paper do not require them.

Moreover, we formalise as locales the notions of a commutative monad, where bind satisfies $m_1 \ggg (\lambda x. m_2 \ggg f x) = m_2 \ggg (\lambda y. m_1 \ggg (\lambda x. f x y))$, and of a discardable monad, where the law $m \ggg m' = m'$ makes it possible to drop a computation whose result is not used.

3 Implementations of Monads and Monad Transformers

In the previous section, we specified the properties of monadic operations abstractly. Now, we provide monad implementations that satisfy these specifications. Some effects are implemented as monad transformers [15,21], which allow us to compose implementations of different effects almost as modularly as the locales specifying them abstractly. In particular, we analyse whether the transformers preserve the specifications of the other effects. All our implementations are polymorphic in the values such that they can be used with any value type, although by the value-monomorphism restriction, each usage must individually commit to one value type.

3.1 The Identity Monad

The simplest monad implementation in our library is the identity monad ident , which models the absence of all effects. It is not really useful in itself, but will be an important building block when combining monads using transformers. The datatype $\alpha \text{ ident}$ is a copy of α with constructor Ident and selector run-ident . To distinguish the abstract monad operations from their implementations, we subscript the latter with the implementation type. The lemma states that $\text{return}_{\text{ident}}$ and $\text{bind}_{\text{ident}}$ satisfy the assumption of the locale monad. Additionally, the identity monad is commutative and discardable.

```
datatype  $\alpha \text{ ident} = \text{Ident} (\text{run-ident}: \alpha)$ 
definition  $\text{return}_{\text{ident}} :: \alpha \Rightarrow \alpha \text{ ident}$  where  $\text{return}_{\text{ident}} = \text{Ident}$ 
definition  $\text{bind}_{\text{ident}} :: \alpha \text{ ident} \Rightarrow (\alpha \Rightarrow \alpha \text{ ident}) \Rightarrow \alpha \text{ ident}$  where
   $m \ggg_{\text{ident}} f = f (\text{run-ident } m)$ 
```

```
lemma monad  $\text{return}_{\text{ident}}$   $\text{bind}_{\text{ident}}$ 
```


3.2 The Probability Monad

The probability monad $\alpha \text{ prob}$ is another basic building block. We use discrete probability distributions [7] and Giry’s probability monad operations `dirac` and `bindpmf`, which we already used in the abstract specification in §2.3. Then, probabilistic choice $\mathfrak{C}_{\text{prob}}$ is just monadic sequencing on $\alpha \text{ pmf}$. The probability monad is commutative and discardable.

```

type-synonym  $\alpha \text{ prob} = \alpha \text{ pmf}$ 
definition returnprob ::  $\alpha \Rightarrow \alpha \text{ prob}$  where returnprob = dirac
definition bindprob ::  $\alpha \text{ prob} \Rightarrow (\alpha \Rightarrow \alpha \text{ prob}) \Rightarrow \alpha \text{ prob}$  where bindprob = bindpmf
definition \mathfrak{C}_{\text{prob}} ::  $\pi \text{ pmf} \Rightarrow (\pi \Rightarrow \alpha \text{ prob}) \Rightarrow \alpha \text{ prob}$  where \mathfrak{C}_{\text{prob}} = bindpmf
lemma monad-prob returnprob bindprob \mathfrak{C}_{\text{prob}}

```

3.3 The Failure and Exception Monad Transformer

Failures and exception handling are implemented as a monad transformer. Thus, these effects can be added to any monad τ . In the value-polymorphic setting, the failure monad transformer takes a monad τ and defines a type constructor `failT` such that $\beta \text{ failT}$ is isomorphic to $(\beta \text{ option}) \tau$. That is, the transformer specialises the value type α of the inner monad to $\beta \text{ option}$. In our value-monomorphic setting, the type variable μ represents the application of τ to the value type, i.e., $\beta \text{ option}$. So, $\mu \text{ failT}$ is just a copy of μ :

```

datatype  $\mu \text{ failT} = \text{FailT}$  (run-fail:  $\mu$ )

```

As `failT`’s operations depend on the inner monad, we fix abstract operations `return` and `bind` in an unnamed context and define `failT`’s operations in terms of them. The line on the left indicates the scope of the context. At the end, which is marked by \perp , the fixed operations become additional arguments of the defined functions. Values in the inner monad now have type $\alpha \text{ option}$. The definitions themselves are standard [21].

```

context fixes return ::  $\alpha \text{ option} \Rightarrow \mu$  and bind ::  $\mu \Rightarrow (\alpha \text{ option} \Rightarrow \mu) \Rightarrow \mu$ 
|
| definition returnfailT ::  $\alpha \Rightarrow \mu \text{ failT}$  where
|   returnfailT  $x = \text{FailT}$  (return (Some  $x$ ))
|
| definition bindfailT ::  $\mu \text{ failT} \Rightarrow (\alpha \Rightarrow \mu \text{ failT}) \Rightarrow \mu \text{ failT}$  where
|   bindfailT  $m \gg=_{\text{failT}}$   $f = \text{FailT}$  (run-fail  $m \gg=$ 
|     ( $\lambda x.$  case  $x$  of None  $\Rightarrow$  return None | Some  $y \Rightarrow$  run-fail ( $f y$ )))
|
| definition failfailT ::  $\mu \text{ failT}$  where failfailT = FailT (return None)
|
| definition catchfailT ::  $\mu \text{ failT} \Rightarrow \mu \text{ failT} \Rightarrow \mu \text{ failT}$  where
|   catchfailT  $m_1 m_2 = \text{FailT}$  (run-fail  $m_1 \gg=$ 
|     ( $\lambda x.$  case  $x$  of None  $\Rightarrow$  run-fail  $m_2$  | Some  $_ \Rightarrow$  return  $x$ ))

```

If `return` and `bind` form a monad, so do `returnfailT` and `bindfailT`, and `failfailT` and `catchfailT` satisfy the effect specification from §2.1, too. The next lemma expresses this.

```

| lemma monad-catch returnfailT bindfailT failfailT catchfailT
|   if monad return bind

```

Clearly, we want to keep using the existing effects of the inner monad. So, we must lift their operations to `failT` and prove that their specifications are preserved. The

lifting is not hard; the continuations of the operations are transformed in the same way as $\text{bind}_{\text{failT}}$ does. Here, we only show how to lift the state operations, where the locale `monad-catch-state` extends `monad-catch` and `monad-state` with `CATCH-GET` and `CATCH-PUT`. Moreover, `failT` also lifts `ϕ`, `alt`, `ask`, `tell`, and `pause`, preserving their specifications. It is commutative if the inner monad is commutative and discardable.

```

context fixes get :: (σ ⇒ μ) ⇒ μ and put :: σ ⇒ μ ⇒ μ
  definition get_failT :: (σ ⇒ μ failT) ⇒ μ failT where
    get_failT f = FailT (get (λs. run-fail (f s)))
  definition put_failT :: σ ⇒ μ failT ⇒ μ failT where
    put_failT s m = FailT (put s (run-fail m))
  lemma monad-catch-state return_failT bind_failT fail_failT catch_failT get_failT put_failT
    if monad-state return bind get put

```

From now on, as the context scope has ended, $\text{return}_{\text{failT}}$ and $\text{bind}_{\text{failT}}$ take the inner monad's operations `return` and `bind` as additional arguments. For example, we obtain a plain failure monad by applying `failT` to `ident`. Interpreting the locale `monad-fail` for $\text{return}_F = \text{return}_{\text{failT}} \text{return}_{\text{ident}}$ and $\text{bind}_F = \text{bind}_{\text{failT}} \text{return}_{\text{ident}} \text{bind}_{\text{ident}}$ and $\text{fail}_F = \text{fail}_{\text{failT}} \text{return}_{\text{ident}}$ yields an executable version of the interpreter `eval` from §2.1, which we refer to as eval_F . Then, Isabelle's code generator and term rewriter both evaluate

$$\text{eval}_F (\lambda x. \text{return}_F (((\lambda _. 0)(x_0 := 5)) x)) (\text{Var } x_0 \oplus \text{Const } 7)$$

to `FailT (Ident (Some 12))`. Given some variable environment $Y :: v \Rightarrow \text{int}$,² we obtain a textbook-style interpreter [22, §3.1.2] as `run-ident (run-fail (evalF (returnF ◦ Y) e))`.

3.4 The State Monad Transformer

The state monad transformer adds the effects of a state monad to some inner monad. The formalisation follows the same ideas as for `failT`, so we only mention the important points. The state monad transformer transforms a monad $\alpha \tau$ into the type $\sigma \Rightarrow (\alpha \times \sigma) \tau$ where σ is the type of states. So, in HOL, the type of values of the inner monad becomes $\alpha \times \sigma$ and μ represents $(\alpha \times \sigma) \tau$.

```
datatype (σ, μ) stateT = StateT (run-state: σ ⇒ μ)
```

Like for `failT`, the state monad operations $\text{return}_{\text{stateT}}$ and $\text{bind}_{\text{stateT}}$ depend on inner monad operations `return` and `bind`. With $\text{get}_{\text{stateT}}$ and $\text{put}_{\text{stateT}}$ defined in the obvious way, the transformer satisfies the specification `monad-state` for state monads.

```

context fixes return :: α × σ ⇒ μ and bind :: μ ⇒ (α × σ ⇒ μ) ⇒ μ
  definition return_stateT :: α ⇒ (σ, μ) stateT where
    return_stateT x = StateT (λs. return (x, s))
  definition bind_stateT :: (σ, μ) stateT ⇒ (α ⇒ (σ, μ) stateT) ⇒ (σ, μ) stateT where
    m >>=_{stateT} f = StateT (λs. run-state f s >>= (λ(x, s'). run-state (f x) s'))
  definition get_stateT :: (σ ⇒ (σ, μ) stateT) ⇒ (σ, μ) stateT where
    get_stateT f = StateT (λs. run-state (f s) s)

```

² Such environments can be nicely handled by applying a reader monad transformer on top (§4).

```

definition putstateT ::  $\sigma \Rightarrow (\sigma, \mu)$  stateT  $\Rightarrow (\sigma, \mu)$  stateT where
  putstateT s m = StateT ( $\lambda \_.$  run-state m s)
lemma monad-state returnstateT bindstateT getstateT putstateT
  if monad return bind

```

The state monad transformer lifts the other effect operations fail, \clubsuit , ask, tell, alt, and pause according to their specifications. But catch cannot be lifted through stateT such that CATCH-GET and CATCH-PUT from §2.4 hold. As our exceptions carry no information, the inner monad cannot pass the state updates before the failure to the handler.

3.5 Composing Monads with Transformers

Composing the two monad transformers failT and stateT with the monad prob, we can now instantiate the probabilistic interpreter from §2.4. As is well known, the order of composition matters. If we first apply failT to prob and then stateT (SFP for short), the resulting interpreter $\text{eval}_{\text{SFP}} E e :: (v \rightarrow \text{int}, (\text{int} \times (v \rightarrow \text{int}))) \text{ option prob failT stateT}$ nests the result state of type $v \rightarrow \text{int}$ inside the option type for failures, i.e., failures do not return a new state. Thus, failures erase state updates, i.e., $\text{put}_{\text{SFP}} s \text{ fail}_{\text{SFP}} = \text{fail}_{\text{SFP}}$, and lazy and eager sampling are equivalent (LAZY-EAGER). Conversely, if we apply failT after stateT to prob (FSP for short), then $\text{eval}_{\text{FSP}} E e :: (v \rightarrow \text{int}, (\text{int option} \times (v \rightarrow \text{int}))) \text{ prob stateT failT}$ and failures do return a new state as only the result type int sits inside option. In particular, $\text{put}_{\text{SFP}} s \text{ fail}_{\text{FSP}} \neq \text{fail}_{\text{FSP}}$ in general, and lazy and eager sampling are not equivalent. We will consider the SFP case further in §4.

3.6 Further Monads and Monad Transformers

Apart from the monad implementations presented so far, our library provides implementations also for the other types of effects mentioned in §2.5. In particular, non-deterministic choice is implemented as a monad transformer based on finite multisets, which works only for commutative inner monads. Moreover, we define a reader (readT) and a writer (writerT) monad transformer. The reader monad transformer differs from stateT only in that no updates are possible. Thus, (ρ, μ) readT leaves the type of values of the inner monad unchanged, as no new state must be returned.

```

datatype ( $\rho, \mu$ ) readT = ReadT (run-read:  $\rho \Rightarrow \mu$ )
context fixes return ::  $\alpha \Rightarrow \mu$  and bind ::  $\mu \Rightarrow (\alpha \Rightarrow \mu) \Rightarrow \mu$ 
definition returnreadT ::  $\alpha \Rightarrow (\rho, \mu)$  readT where
  returnreadT x = ReadT ( $\lambda \_.$  return x)
definition bindreadT :: ( $\rho, \mu$ ) readT  $\Rightarrow (\alpha \Rightarrow (\rho, \mu)$  readT)  $\Rightarrow (\rho, \mu)$  readT where
  m  $\gg_{\text{readT}}$  f = ReadT ( $\lambda r.$  run-read m r  $\gg_{\text{readT}}$  ( $\lambda x.$  run-read (f x) r))
definition askreadT :: ( $\rho \Rightarrow (\rho, \mu)$  readT)  $\Rightarrow (\rho, \mu)$  readT where
  askreadT f = ReadT ( $\lambda r.$  run-read (f r) r)
definition failreadT :: ( $\mu \Rightarrow (\rho, \mu)$  readT) where failreadT fail = ReadT ( $\lambda \_.$  fail)

```

Resumptions are formalised as a plain monad using the codatatype
codatatype (o, ι, α) resumption = Done α | Pause o ($\iota \Rightarrow (o, \iota, \alpha)$ resumption)

Unfortunately, we cannot define resumptions as a monad transformer in HOL despite the restriction to monomorphic values. The reason is that for a transformer with inner monad τ , the second argument of the constructor `Pause` would have to be of type $\iota \Rightarrow (o, \iota, \alpha)$ resumption τ , i.e., the codatatype would recurse through the unspecified type constructor τ . This is not supported by Isabelle’s codatatype package [2] and, in fact, for some choices of τ , e.g., unbounded nondeterminism, the resumption transformer type does not exist in HOL at all. For the same reason, we cannot have other monad transformers that have similar recursive implementation types. Therefore, we fail to modularly construct all combinations of effects. For example, probabilistic resumptions with failures [17] are out of reach and must still be constructed from scratch.

3.7 Overloading the Monad Operations

When several monad transformers are composed, the monad operations quickly become large HOL terms as the transformer’s operations take the inner monad’s as explicit arguments. These large terms must be handled by the inference kernel, the type checker, the parser, and the pretty-printer, even if locale interpretations hide them from the user using abbreviations. To improve readability and the processing time of Isabelle, our library also defines the operations as single constants which are overloaded for the different monad implementations using recursion on types [26]. As overloading does not need these explicit arguments, it thus avoids the processing times for unification, type checking, and (un)folding of abbreviations. Yet, Isabelle’s check against cyclic definitions [13] fails to see that the resulting dependencies must be acyclic (as the inner monad is always a type argument of the outer monad). So, we moved these overloaded definitions to a separate file and marked them as unchecked.³ Overloading is just a syntactic convenience, on which the library and the examples in this paper do not rely. If users want to use it, they are responsible for not exploiting these unchecked dependencies.

4 Moving Between Monad Instances

Once all variables have been sampled eagerly, the evaluation of the expression itself is deterministic. Thus, the actual evaluation need not be done in a monad as complex as FSP or SFP. It suffices to work in a reader-failure monad with operations `fail` and `ask`, which we obtain by applying the monad transformers `readT` and `failT` to `ident` (RFI for short). Such simpler monads have the advantage that reasoning becomes easier as more laws hold. We now explain how the theory of representation independence [20] can be used to move between different monad instances by going from SFP to RFI. This ultimately yields a theorem that characterises `evalSFP` in terms of `evalRFI`. So, in general, this approach makes it possible to switch in the middle of a bigger proof from a complicated monad to a much simpler one.

³ Isabelle’s `adhoc-overloading` feature, which resolves overloading during type checking, cannot be used either as it does not support recursive resolutions. For example, resolving `return :: $\alpha \Rightarrow \alpha$ option ident failT` takes two steps: first to `returnfailT return` and then to `returnfailT returnident`. The second step fails due to the intricate interleaving of type checking and resolution. Even if this is just a implementation issue, resolving overloading during type checking prevents definitions that are generic in the monad, which general overloading supports.

Let us first deal with sampling. To go from α prob to β ident, we use a relation $\mathbb{I}\mathbb{P}(A)$ between α ident and β prob since relations work better with higher-order functions than equations. Following Huffman and Kunčar [10], we call such relations correspondence relations. It is parametrised by a relation A between the values, which we will use later to express the differences in the values due to the monad transformers changing the value type of the inner monad. In detail, $\mathbb{I}\mathbb{P}(A)$ relates a value $\text{Ident } x$ to the one-point distribution $\text{dirac } y$ iff A relates x to y . Then, the monad operations of ident and prob respect this relation. Respectfulness is formalised using the function relator $A \Rightarrow B$ defined by $(f, g) \in A \Rightarrow B$ iff $(x, y) \in A$ implies $(f(x), g(y)) \in B$ for all x and y . Then, the monad operations respecting $\mathbb{I}\mathbb{P}(A)$ is expressed by the following two conditions:

- $(\text{return}_{\text{ident}}, \text{return}_{\text{prob}}) \in A \Rightarrow \mathbb{I}\mathbb{P}(A)$ and
- $(\text{bind}_{\text{ident}}, \text{bind}_{\text{prob}}) \in \mathbb{I}\mathbb{P}(A) \Rightarrow (A \Rightarrow \mathbb{I}\mathbb{P}(A)) \Rightarrow \mathbb{I}\mathbb{P}(A)$.

Note the similarity between the relations and the types of the monad operations, where A and $\mathbb{I}\mathbb{P}$ take the roles of the type variables for values and of the monad type constructor, respectively. As the monad transformers failT and stateT are relationally parametric in the inner monad and eval is parametric in the monad, we prove the following relation between the evaluators automatically using Isabelle/HOL's Transfer prover [10]

$$(\text{eval}_{\text{SFP}} \text{lookup}_{\text{SFP}} e, \text{eval}_{\text{SFI}} \text{lookup}_{\text{SFI}} e) \in \text{rel}_{\text{stateT}} (\text{rel}_{\text{failT}} (\mathbb{I}\mathbb{P}(=))) \quad (4)$$

where SFI refers to the state-failure-identity composition of monads, $(=)$ is the identity relation, and $\text{rel}_{\text{stateT}}$ and $\text{rel}_{\text{failT}}$ are the relators for the datatypes stateT and failT [2]. Formally, the relators lift relations on the inner monad to relations on the transformed monad. For example, $(m_1, m_2) \in \text{rel}_{\text{stateT}} M$ iff $(\text{run-state } m_1 \ s, \text{run-state } m_2 \ s) \in M$ for all s , and $(m_1, m_2) \in \text{rel}_{\text{failT}} M$ iff $(\text{run-fail } m_1, \text{run-fail } m_2) \in M$. Intuitively, (4) states that in the monads SFP and SFI, eval behaves the same with respect to states updates and failure and the results are the same; in particular, the evaluation is deterministic.

In the following, we use the property of a relator rel that if M is the graph $\text{Gr } f$ of a function f , then $\text{rel } M$ is the graph of the function map f , where map is the canonical map function for the relator. For example, $\text{map}_{\text{failT}} f = \text{FailT} \circ f \circ \text{run-fail}$, so

$$\text{rel}_{\text{failT}} (\text{Gr } f) = \text{Gr } (\text{map}_{\text{failT}} f) \quad (5)$$

where $(x, y) \in \text{Gr } f$ iff $f \ x = y$. Isabelle's datatype package automatically proves these relator-graph identities. The correspondence relation $\mathbb{I}\mathbb{P}$ satisfies a similar law: $\mathbb{I}\mathbb{P}(\text{Gr } f) = \text{Gr } (\text{map}_{\mathbb{I}\mathbb{P}} f)$ where $\text{map}_{\mathbb{I}\mathbb{P}} f = \text{dirac} \circ f \circ \text{run-ident}$.

Having eliminated probabilities, we next switch from the state monad transformer to the reader monad transformer. We again define a correspondence relation $\mathbb{R}\mathbb{S}(s, M)$ between readT and stateT . It takes as parameters the environment s and the correspondence relation M between the inner monads. It relates the two monadic values m_1 and m_2 iff M relates the results of running m_1 and m_2 on s , i.e., $(\text{run-read } m_1 \ s, \text{run-state } m_2 \ s) \in M$. Again, we show that the monad operations respect $\mathbb{R}\mathbb{S}(s, M)$ as formalised below. As readT and stateT are monad transformers, we assume that the operations of the inner monads respect M . These assumptions can be expressed using \Rightarrow since the inner operations are arguments to readT 's and stateT 's operations. Here, $A \triangleleft \times s$ adapts the relation A on values to stateT 's change of the value type from α to $\alpha \times \sigma$; $(x, (y, s')) \in A \triangleleft \times s$ iff $(x, y) \in A$ and $s' = s$, i.e., A relates the results and the state is not updated.

- $(\text{return}_{\text{readT}}, \text{return}_{\text{stateT}}) \in (A \ltimes s \Rightarrow M) \Rightarrow A \Rightarrow \mathbb{RS}(s, M)$,
- $(\text{bind}_{\text{readT}}, \text{bind}_{\text{stateT}}) \in$
 $(M \Rightarrow (A \ltimes s \Rightarrow M) \Rightarrow M) \Rightarrow \mathbb{RS}(s, M) \Rightarrow (A \Rightarrow \mathbb{RS}(s, M)) \Rightarrow \mathbb{RS}(s, M)$,
- $(\text{ask}_{\text{readT}}, \text{get}_{\text{stateT}}) \in (\{(s, s)\} \Rightarrow \mathbb{RS}(s, M)) \Rightarrow \mathbb{RS}(s, M)$, and
- $(\text{fail}_{\text{readT}}, \text{fail}_{\text{stateT}}) \in M \Rightarrow \mathbb{RS}(s, M)$,

Then, by representation independence, the Transfer package automatically proves the following relation between eval_{RFI} and eval_{SFI} , where $\text{lookup}_{\text{RFI}}$ uses $\text{ask}_{\text{readT}}$ instead of $\text{get}_{\text{stateT}}$, and $\text{rel}_{\text{ident}}$ and $\text{rel}_{\text{option}}$ are the relators for the datatypes `ident` and `option`.

$$(\text{eval}_{\text{RFI}} \text{lookup}_{\text{RFI}} e, \text{eval}_{\text{SFI}} \text{lookup}_{\text{SFI}} e) \in \mathbb{RS}(s, \text{rel}_{\text{failT}} (\text{rel}_{\text{ident}} (\text{rel}_{\text{option}} (= \ltimes s))))$$

This says that running `eval` in RFI and SFI computes the same result, has the same behaviour with respect to state queries and failures, and does not update the state.

Actually, we can go from SFP directly to RFI, without the monad SFI as a stepping stone, thanks to \mathbb{IP} taking a relation on the value types:

$$(\text{eval}_{\text{RFI}} \text{lookup}_{\text{RFI}} e, \text{eval}_{\text{SFP}} \text{lookup}_{\text{SFP}} e) \in \mathbb{RS}(s, \text{rel}_{\text{failT}} (\mathbb{IP}(\text{rel}_{\text{option}} (= \ltimes s)))) \quad (6)$$

As $= \ltimes s$ is the graph of $\lambda a. (a, s)$, using only the graph properties like (5) of \mathbb{IP} and the relators, and using \mathbb{RS} 's definition, we derive the characterisation of eval_{SFP} from (6):

$$\text{run-state } (\text{eval}_{\text{SFP}} \text{lookup}_{\text{SFP}} e) s =$$

$$\text{map}_{\text{failT}} (\text{map}_{\mathbb{IP}} (\text{map}_{\text{option}} (\lambda a. (a, s)))) (\text{run-read } (\text{eval}_{\text{RFI}} \text{lookup}_{\text{RFI}} e) s)$$

where $\text{map}_{\text{failT}}$ and $\text{map}_{\text{option}}$ are the canonical map functions for `failT` and `option`. Thus, instead of reasoning about eval_{SFP} in SFP, we can conduct our proofs in the simpler monad RFI. For example, as RFI is commutative, subexpressions can be evaluated in any order. Thus, we get the following identity expressing the reversed evaluation order (and a similar one for \odot).⁴

$$\text{eval}_{\text{RFI}} E (e_1 \oplus e_2) = \text{eval}_{\text{RFI}} E e_2 \gg_{\text{RFI}} (\lambda j. \text{eval}_{\text{RFI}} E e_1 \gg_{\text{RFI}} (\lambda i. \text{return}_{\text{RFI}} (i + j)))$$

In summary, we have demonstrated a generic approach to switch from a complicated monad to a much simpler one. Conceptually, the correspondence relations \mathbb{IP} and \mathbb{RS} just embed one monad or monad transformer (`ident` and `readT`) in a richer one (`prob` and `stateT`). It is precisely this embedding that ultimately yields the map functions in the characterisation. In this functional view, the respectfulness conditions express that the embedding is a monad homomorphism. Yet, we use relations for the embedding instead of functions because only relations work for higher-order operations in a compositional way.

The reader may wonder why we go through all the trouble of defining correspondence relations and showing respectfulness and parametricity. Indeed, in this example, it would probably have been easier to simply perform an induction over expressions and prove the equation directly. The advantage of our approach is that it does not rely on the concrete definition of `eval`. It suffices to know that `eval` is parametric in the monad, which Isabelle derives automatically from the definition. This automated approach therefore scales to arbitrarily complicated monadic functions whereas induction proofs do not. Moreover, note that the correspondence relations and respectfulness lemmas only depend on the monads. They can therefore be reused for other monadic functions.

⁴ Following the “as abstract as possible” spirit of this paper, we actually proved the identities in the locale of commutative monads and showed that `readT` is commutative if its inner monad is.

5 Related work

Huffman et al. [11,9] formalise the concept of value-polymorphic monads and several monad transformers in Isabelle/HOLCF, the domain theory library of Isabelle/HOL. They circumvent HOL’s type system restrictions by projecting everything into HOLCF’s universal domain of computable values. That is, they trade in HOL’s set-theoretic model with its simple reasoning rules for a domain-theoretic model with ubiquitous \perp values and strictness side conditions. This way, they can define a resumption monad transformer (for computable continuations). Being tied to domain theory, their library cannot be used to model effects of *plain* HOL functions, which is our goal, the strictness assumptions make their laws and proofs more complicated than ours, and functions defined with HOLCF do not work with Isabelle’s code generator. Still, their idea of projecting everything into a universal type could also be adapted to plain HOL, albeit only for a restricted class of monads; achieving a similar level of automation and modularity would require a lot more effort than our approach, which uses only existing features of Isabelle.

Gibbons and Hinze [6] axiomatize monads and effects using Haskell-style type constructor classes and use the algebraic specification to prove identities between Haskell programs, similar to our abstract locales in §2. Their specification of state effects omits GET-CONST, but they later assume that it holds [6, §10.2]. Being value-polymorphic, their operations do not need our continuations and the laws are therefore simpler. In particular, no new assumptions are typically needed when monad specifications are combined. In contrast, our continuations sometimes require interaction assumptions like SAMPLE-GET. Gibbons and Hinze only consider reasoning in the abstract setting and do not discuss the transition to concrete implementations and the relations between implementations. Also, they do not prove that monad implementations satisfy their specifications. Later, Jeuring et al. [12] showed that the implementations in Haskell do not satisfy them because of strictness issues similar to the ones in Huffman’s work.

Lobo Vesga [16] formalised some of Gibbons’ and Hinze’s examples in Agda. She does not need assumptions for the continuations like we do as value-polymorphic monads can be directly expressed in Agda. Like Gibbons and Hinze, she does not study the connection between specifications and implementations. Thanks to the good proof automation in Isabelle, our mechanised proofs are much shorter than hers, which are as detailed as Gibbons’ and Hinze’s pen-and-paper proofs.

Lochbihler and Schneider [19] implemented support for equational reasoning about applicative functors, which are more general than monads. They focus on lifting identities on values to a concrete applicative functor. Reasoning with abstract applicative functors is not supported. Like monads, the concept of an applicative functor cannot be expressed as a predicate in HOL. Moreover, the applicative operations do not admit value monomorphisation like monads do, as the type of \diamond contains applications of the functor type constructor τ to $\alpha \Rightarrow \beta$, α , and β . So, monads seem to be the right choice, even though we could have defined the interpreter eval applicatively (but not, e.g., memoisation).

6 Conclusion

We have presented a library of abstract monadic effect specifications and their implementations as monads and monad transformers in Isabelle/HOL. We illustrated its usage and

the elegance of reasoning using a monadic interpreter. The type system of HOL forced us to restrict the monads to monomorphic values. Monomorphic values work well when the reasoning involves only a few monadic functions like in our running example. In larger projects, this restriction can become a limiting factor. Nevertheless, in our project on formalising computational soundness results,⁵ we successfully formalised and reasoned about several complicated serialisers and parsers for symbolic messages of security protocols. In that work, reasoning abstractly about effects and being able to move from one monad instance to another were crucial. More concretely, the serialiser converts symbolic protocol messages into bitstrings. The challenges were similar to those of our interpreter `eval`. Serialisation may fail when the symbolic message is not well-formed, similar to division by zero in the interpreter. When serialisation encounters a new nonce, it randomly samples a fresh bitstring, which must also be used for serialising further occurrences of the same nonce. We formalised this similar to the memoisation of variable evaluation in the interpreter. A further challenge not present in the interpreter was that the serialiser must also record the serialisation of all subexpressions such that the parser can map bitstrings generated by the serialiser back to symbolic messages without calling a decryption oracle or inverting a cryptographic hash function. The construction relied on the invariant that the recorded values were indeed generated by the serialiser, but such an invariant cannot be expressed easily for a probabilistic, stateful function. We therefore formalised also the switch from lazy to eager sampling for the serialiser (lazy sampling was needed to push the randomisation of encryptions into an encryption oracle) and the switch to a read-only version without recording of results using similar to our example in §4.

Instead of specifying effects abstractly and composing them using monad transformers, we obviously could have formalised everything in a sufficiently rich monad that covers all the effects of interest, e.g., continuations. Then, there would be no need for abstract specifications as we could work directly with a concrete monad as usual, where our reasoning on the abstract level could be mimicked. But we would deprive ourselves of the option of going to a specific monad that covers precisely the effects needed. Such specialisation has two advantages: First, as shown in §4, simpler monads satisfy more laws, e.g., commutativity, which make the proofs easier. Second, concrete monads can have dedicated setups for reasoning and proof automation that are not available in the abstract setting. Our library achieves the best of both worlds. We can reason abstractly and thus achieve generality. When this gets too cumbersome or impossible, we can switch to a concrete monad, continuing to use the abstract properties already proven.

In the long run, we can imagine a definitional package for monads and monad transformers that composes concrete value-polymorphic monad transformers. Similar to how Isabelle’s `datatype` package composes bounded natural functors [2], such a package must perform the construction and the derivation of all laws afresh for every concrete combination of monads, as value-polymorphic monads lie beyond HOL’s expressiveness. When combined with a reinterpretation framework for theories, we could model effects and reason about them abstractly and concretely without the restriction to monomorphic values.

Acknowledgements We thank Dmitriy Traytel and the anonymous reviewers for suggesting many improvements to the presentation. This work is supported by the Swiss National Science Foundation grant 153217 “Formalising Computational Soundness for Protocol Implementations”.

⁵ <http://www.infsec.ethz.ch/research/projects/FCSPI.html>

References

1. Ballarin, C.: Locales: A module system for mathematical theories. *J. Automat. Reason.* 52(2), 123–153 (2014)
2. Blanchette, J.C., Hölzl, J., Lochbihler, A., Panny, L., Popescu, A., Traytel, D.: Truly modular (co)datatypes for Isabelle/HOL. In: ITP 2014. LNCS, vol. 8558, pp. 93–110. Springer (2014)
3. Bulwahn, L., Krauss, A., Haftmann, F., Erkök, L., Matthews, J.: Imperative functional programming with Isabelle/HOL. In: TPHOLs 2008. LNCS, vol. 5170, pp. 134–149. Springer (2008)
4. Eberl, M., Hölzl, J., Nipkow, T.: A verified compiler for probability density functions. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 80–104. Springer (2015)
5. Erwig, M., Kollmansberger, S.: Functional pearls: Probabilistic functional programming in Haskell. *Journal of Functional Programming* 16, 21–34 (2006)
6. Gibbons, J., Hinze, R.: Just do it: Simple monadic equational reasoning. In: ICFP 2011. pp. 2–14. ACM (2011)
7. Hölzl, J., Lochbihler, A., Traytel, D.: A formalized hierarchy of probabilistic system types. In: ITP 2015. LNCS, vol. 9236, pp. 203–220. Springer (2015)
8. Homeier, P.V.: The HOL-Omega logic. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 244–259. Springer (2009)
9. Huffman, B.: Formal verification of monad transformers. In: ICFP 2012. pp. 15–16. ACM (2012)
10. Huffman, B., Kunčar, O.: Lifting and Transfer: A modular design for quotients in Isabelle/HOL. In: CPP 2013. LNCS, vol. 8307, pp. 131–146. Springer (2013)
11. Huffman, B., Matthews, J., White, P.: Axiomatic constructor classes in Isabelle/HOLCF. In: Hurd, J., Melham, T. (eds.) TPHOLs 2005. LNCS, vol. 3603, pp. 147–162. Springer (2005)
12. Jeuring, J., Jansson, P., Amaral, C.: Testing type class laws. In: Haskell 2012. pp. 49–60. ACM (2012)
13. Kunčar, O.: Correctness of Isabelle’s cyclicity checker: Implementability of overloading in proof assistants. In: CPP 2015. pp. 85–94. ACM (2015)
14. Lammich, P., Tuerk, T.: Applying data refinement for monadic programs to Hopcroft’s algorithm. In: ITP 2012. LNCS, vol. 7406, pp. 166–182. Springer (2012)
15. Liang, S., Hudak, P., Jones, M.: Monad transformers and modular interpreters. In: POPL 1995. pp. 333–343. ACM (1995)
16. Lobo Vesga, E.: Hacia la formalización del razonamiento ecuacional sobre mónadas. Tech. rep., Universidad EAFIT (2013), <http://hdl.handle.net/10784/4554>
17. Lochbihler, A.: Probabilistic functions and cryptographic oracles in higher order logic. In: Thiemann, P. (ed.) ESOP 2016. LNCS, vol. 9632, pp. 503–531. Springer (2016)
18. Lochbihler, A.: Effect polymorphism in higher-order logic. *Archive of Formal Proofs* (2017), http://isa-afp.org/entries/Monomorphic_Monad.shtml, Formal proof development
19. Lochbihler, A., Schneider, J.: Equational reasoning with applicative functors. In: Blanchette, J.C., Merz, S. (eds.) ITP 2016. LNCS, vol. 9807, pp. 252–273. Springer (2016)
20. Mitchell, J.C.: Representation independence and data abstraction. In: POPL 1986. pp. 263–276. ACM (1986)
21. Moggi, E.: An abstract view of programming languages. Tech. Rep. ECS-LFCS-90-113, LFCS, School of Informatics, University of Edinburgh (1990)
22. Nipkow, T., Klein, G.: *Concrete Semantics*. Springer (2014)
23. Nipkow, T., Paulson, L.C.: Proof pearl: Defining functions over finite sets. In: Hurd, J., Melham, T. (eds.) TPHOLs 2005. LNCS, vol. 3603, pp. 385–396. Springer (2005)
24. Ramsey, N., Pfeffer, A.: Stochastic lambda calculus and monads of probability distributions. In: POPL 2002. pp. 154–165. ACM (2002)
25. Wadler, P.: Monads for functional programming. In: Jeuring, J., Meijer, E. (eds.) *Advanced Functional Programming*, LNCS, vol. 925, pp. 24–52. Springer (1995)
26. Wenzel, M.: Type classes and overloading in higher-order logic. In: Gunter, E.L., Felty, A. (eds.) TPHOLs 1997. LNCS, vol. 1275, pp. 307–322. Springer (1997)