THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Dijkstra Monads for Free

Citation for published version:

Ahman, D, Hritcu, C, Maillard, K, Martínez, G, Plotkin, G, Protzenko, J, Rastogi, A & Swamy, N 2017, Dijkstra Monads for Free. in *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages*. ACM SIGPLAN Notices, no. 1, vol. 52, ACM, pp. 515-529, 44th ACM SIGPLAN Symposium on Principles of Programming Languages 2017, Paris, France, 15/01/17. <https://doi.org/10.1145/3009837.3009878>

Digital Object Identifier (DOI):

[10.1145/3009837.3009878](https://doi.org/10.1145/3009837.3009878)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Dijkstra Monads for Free

Danel Ahman¹ Cătălin Hrițcu² Guido Martínez² Gordon Plotkin¹
Jonathan Protzenko³ Aseem Rastogi³ Nikhil Swamy³

¹University of Edinburgh

²Inria Paris

³Microsoft Research

Abstract

Dijkstra monads are a means by which a dependent type theory can be enhanced with support for reasoning about effectful code. These specification-level monads computing weakest preconditions, and their closely related counterparts, *Hoare monads*, provide the basis on which verification tools like F^* , Hoare Type Theory (HTT), and Ynot are built. In this paper we show that Dijkstra monads can be derived “for free” by applying a continuation-passing style (CPS) translation to the standard monadic definitions of the underlying computational effects.

Automatically deriving Dijkstra monads provides a correct-by-construction and efficient way of reasoning about user-defined effects in dependent type theories. We demonstrate these ideas in EMF^* , a new dependently typed calculus, validating it both by formal proof and via a prototype implementation within F^* . Besides equipping F^* with a more uniform and extensible effect system, EMF^* enables within F^* a mixture of intrinsic and extrinsic proofs that was previously impossible.

1. Introduction

Monads are a versatile concept:

- From Moggi (1989), they are used to give semantics to call-by-value reduction.
- From Moggi (1989), Wadler (1990, 1992), Filinski (1994, 1999, 2010), Benton et al. (2002) and others, they are used as a way to introduce effects into a functional language.
- From Moggi (1989); Flanagan et al. (1993); Wadler (1994) and others, they provide a foundation on which to understand program transformations, notably CPS.

This paper brings together this threefold use of monads to improve upon a fourth use of monads, namely “Dijkstra monads”, a recent proposal by Swamy et al. (2013, 2016) and Jacobs (2015), who suggest using monads of predicate transformers (like Dijkstra’s weakest preconditions), to verify effectful programs within a functional programming language.

1.1 Example: A Dijkstra monad for stateful computations

In Dijkstra’s (1975) weakest precondition semantics, stateful computations transform postconditions relating results and final state, to preconditions on input states. This gives rise to a Dijkstra monad:

$$\text{WP_ST } a = \text{post } a \rightarrow \text{pre} \quad \text{where } \text{post } a = (a * \text{state}) \rightarrow \text{Type} \\ \text{pre} = \text{state} \rightarrow \text{Type}$$

$$\text{return_WP_ST } x \text{ post } s0 = \text{post } (x, s0) \\ \text{bind_WP_ST } f g \text{ post } s0 = f (\lambda (x, s1) \rightarrow g x \text{ post } s1) s0$$

The *weakest precondition (WP)* of a pure term e is computed to be $\text{return_WP_ST } e$, and the WP of the sequential composition

$\text{let } x = e1 \text{ in } e2$ is computed to be $\text{bind_WP_ST } wp1 (\lambda x. wp2)$, where $wp1$ and $wp2$ are the WPs of $e1$ and $e2$ respectively.

Based on such a construction, Swamy et al. (2013, 2016) (building on previous work by Nanevski et al. (2008)) devised a type system to compute WPs for higher-order, effectful programs, enabling their verification within a dependently typed logic. However, these constructions require a meta-theoretic argument to establish their soundness with respect to the semantics of effectful programs. This typically requires proofs of various correctness and admissibility conditions, including the monad laws and monotonicity.

1.2 For fun: Deriving Dijkstra monads

Rather than being given manually, we show that these predicate transformers can be automatically derived by CPS’ing purely functional definitions of monadic effects (with answer type Type). For instance, rather than defining WP_ST , one can simply compute it by CPS’ing the familiar ST monad (i.e., $\text{state} \rightarrow a * \text{state}$), deriving

$$\text{WP_ST } a = ((a * \text{state}) \rightarrow \text{Type}) \rightarrow \text{state} \rightarrow \text{Type} \quad (\text{unfolded})$$

We introduce DM , a simply typed, pure, monadic metalanguage in which one can, in the spirit of Wadler (1992), define a variety of monadic effects, ranging from state and exceptions, to continuations. We define a type-directed CPS translation for this language and show that monads are translated to Dijkstra monads, i.e., monotone, conjunctive, predicate transformer monads.

1.3 For profit: Program verification with user-defined effects

We apply our technique of deriving Dijkstra monads to F^* (Swamy et al. 2016), a dependently typed programming language that already has at its core a system of primitive effects specified using Dijkstra monads. Our goal is to make F^* ’s effect system easier to configure and extensible beyond the primitive effects it already supports; we proceed as follows.

A core dependent type theory with monadic reflection To formally study our improvements to F^* , we define a new dependently typed core calculus, EMF^* (for Explicitly Monadic F^*) that features an extensible effect system. EMF^* is loosely based on the Calculus of Constructions (Coquand and Huet 1988) with (among other features): (1) a predicative hierarchy of non-cumulative universes; (2) a weakest-precondition calculus for pure programs; (3) refinement types; and (4) a facility for representing user-defined effects using the monadic reflection and reification of Filinski (1994), adapted to the dependently typed setting. New effects can be introduced into the language by defining them in terms of the built-in pure constructs, related to each other via monad morphisms, and each such effect obtains a suitable weakest precondition calculus derived from the underlying pure WPs. We prove the calculus strongly normalizing and the WP calculus sound for total correctness verification, for both pure and effectful programs.

Translating DM to EMF* To extend EMF* with a new effect, one starts by defining a monadic effect (say ST) in DM. Via the CPS transformation, we obtain the Dijkstra variant of that effect (WP_{ST}) as a predicate transformer monad in EMF*. A second translation from DM produces expression-level terms to represent monadic computations in EMF*. A logical relations proof shows that monadic computations are correctly specified by their predicate transformers. We show examples of these translations at work for monadic effects including state, exceptions, information-flow control, continuations, and some combinations thereof.

Intrinsic and extrinsic proofs in EMF* Effectful programs in EMF* can be proven correct using one or both of two different reasoning styles. First, using the WP calculus, programs can be proven intrinsically, by decorating their definitions with specifications that must be proven to be at least as strong as their WPs. We refer to this as the *intrinsic* style, already familiar to users of F*, and other tools like HTT (Nanevski et al. 2008), Dafny (Leino 2010), and Why3 (Filliâtre and Paskevich 2013).

Second, through monadic reification, EMF* allows effectful programs to be revealed as their underlying pure implementations. Once reified, one can reason about them via the computational behavior of their definitions. As such, one may define effectful programs with relatively uninformative types, and prove properties about them as needed, via reification. This *extrinsic* style of proving is familiar to users of systems like Coq or Isabelle, where it is routinely employed to reason about pure functions. In EMF*, this style extends smoothly to terminating effectful programs.

Primitive effects in a call-by-value semantics We see EMF* as a meta-language in which to analyze and describe the semantics of terms in an object language, EMF_{ST}*, a call-by-value programming language with primitive state. In the spirit of Moggi (1989), we show that EMF* programs that treat their ST effect abstractly soundly model EMF_{ST}* reductions—technically, we prove a simulation between EMF_{ST}* and EMF*. As such, our work is a strict improvement on the prior support for primitive effects in F*: despite programming and proving programs in a pure setting, stateful programs can still be compiled to run efficiently in the primitively effectful EMF_{ST}*, while programs with other user-defined effects (e.g., information-flow control) can, unlike before, be executed via their pure encodings.

A prototype implementation for F* We have adapted F* to benefit from the theory developed in this paper, using a subset of F* itself as an implementation of DM, and viewing EMF_{ST}* as a model of its existing extraction mechanism to OCaml. Programmers can now configure F*'s effect system using simple monadic definitions, use F* to prove these definitions correct, and then use our CPS transformation to derive the Dijkstra monads required to configure F*'s existing type-checker. To benefit from the new extrinsic proving capabilities, we also extended F* with two new typing rules, and changed its normalizer, to handle monadic reflection and reification.

Several examples show how our work allows F* to be easily extended beyond the primitive effects already supported, without compromising its efficient primitive effect compilation strategy; and how the new extrinsic proof style places effectful reasoning in F* on an equal footing with its support for reasoning about pure programs.

1.4 Summary of contributions

The central contribution of our work is designing three closely related lambda calculi, studying their metatheory and the connections between them, and applying them to provide a formal and practical foundation for a user-extensible effect system for F*. Specifically,

- (1) EMF*: A new dependent type theory with user-extensible, monadic effects; monadic reflection and reification; WPs; and

refinement types. We prove that EMF* is strongly normalizing and that its WPs are sound for total correctness (§3).

- (2) DM: A simply typed language to define the expression-level monads that we use to extend EMF* with effects. We define a CPS transformation of DM terms to derive Dijkstra monads from expression-level monads, as well as an elaboration of DM terms to EMF*. Moreover, elaborated terms are proven to be in relation with their WPs (§4).
- (3) EMF_{ST}*: A call-by-value language with primitive state, whose reductions are simulated by well-typed EMF* terms (§5).
- (4) An implementation of these ideas within F* (§3.5, §4.6) and several examples of free Dijkstra monads for user-defined effects (§2). We highlight, in particular, the new ability to reason extrinsically about effectful terms.

The auxiliary materials for this paper (<https://www.fstar-lang.org/papers/dm4free>) contain appendices with complete definitions and proofs for the formal results in §3, §4 (Appendix A below), and §5. The F* source code (<https://github.com/FStarLang/FStar>) now includes the extensions from §3.5 and §4.6 and the examples from §2 (<https://github.com/FStarLang/FStar/tree/master/examples/dm4free>).

2. Illustrative examples

We illustrate our main ideas using several examples from F*, contrasting with the state of affairs in F* prior to our work. We start by presenting the core WP calculus for pure programs (§2.1), then illustrate how state can be added to it (§2.2). After showing our basic methodology on state (§2.3 and §2.4), we present a few additional examples, including combining exceptions with state (§2.5), information-flow control (§2.6), and continuations (§2.7).

Notation: The syntax $\lambda(b_1) \dots (b_n) \rightarrow t$ introduces a lambda abstraction, where b_i ranges over binding occurrences $x:t$ declaring a variable x at type t . The type $b_1 \rightarrow \dots \rightarrow b_n \rightarrow c$ is the type of a curried function, where c is a computation type—we emphasize the lack of enclosing parentheses on the b_i . We write just the type in b when the name is irrelevant, and $t \rightarrow t'$ for $t \rightarrow \text{Tot } t'$.

2.1 WPs for pure programs

Reasoning about purely functional programs is a relatively well-understood activity: the type theories underlying systems like Coq, Agda, and F* are already well-suited to the task. Consider proving that pure term $\text{sqr} = \lambda(x:\text{int}) \rightarrow x * x$ always returns a non-negative integer. A natural strategy is an *extrinsic* proof, which involves giving sqr a simple type such as $\text{int} \rightarrow \text{Tot int}$, the type of total functions on integers, and then proving a lemma $\forall x. \text{sqr } x \geq 0$. In the case of F*, the proof of the lemma involves, first, a little computation to turn the goal into $\forall x. x * x \geq 0$, and then reasoning in the theory of integer arithmetic of the Z3 SMT solver (de Moura and Bjørner 2008) to discharge the proof.

An alternative *intrinsic* proof style in F* involves giving sqr type $x:\text{int} \rightarrow \text{Pure int} (\lambda \text{post} \rightarrow \forall y. y \geq 0 \implies \text{post } y)$, a dependent function type of the form $x:t \rightarrow c$, where the formal parameter $x:t$ is in scope in the *computation type* c to the right of the arrow. Computation types c are either $\text{Tot } t$ (for some type t) or of the form $M \ t \ \text{wp}$, where M is an effect label, t is the result type of the computation, and wp is a predicate transformer specifying the semantics of the computation. The computation type we give to sqr is of the form $\text{Pure } t \ \text{wp}$, the type of t -returning pure computations described by the predicate transformer wp : $(t \rightarrow \text{Type}) \rightarrow \text{Type}$, a function taking postconditions on the result (predicates of type $t \rightarrow \text{Type}$), to preconditions. These predicate transformers form a Dijkstra monad. In this case, the wp states that to prove any property post of $\text{sqr } x$, it suffices to prove

post y , for all non-negative y —as such, it states our goal that $\text{sqr } x$ is non-negative. To prove sqr can be given this type, F^* infers a weakest precondition for $\text{sqr } x$, namely $\lambda \text{ post} \rightarrow \text{post } (x * x)$ and aims to prove that the predicate transformer we specified is at least as strong as the weakest one it inferred: $\forall \text{ post}. (\forall y. y \geq 0 \implies \text{post } y) \implies \text{post } (x * x)$, which is discharged automatically by Z3. For pure programs, this intrinsic proof style may seem like overkill and, indeed, it often is. But, as we will see, this mechanism for reasoning about pure terms via WPs is a basic capability which we can leverage for reasoning about terms with more complex, effectful semantics.

2.2 Adding WPs for state

Consider proving that `incr _ = let x = get() in put (x + 1)` produces an output state greater than its input state. Since this program has the state effect, a proof by extrinsic reasoning is not completely straightforward, because reducing an effectful computation within a logic may not be meaningful. Instead, tools like Ynot (Chlipala et al. 2009), HTT (Nanevski et al. 2008), and F^* only support the intrinsic proof style. In the case of F^* , this involves the use of a computation type $ST' t \text{ wp}$, where $\text{wp} : WP_ST t$ and for our simple example we take $WP_ST t = ((t * \text{int}) \rightarrow \text{Type}) \rightarrow \text{int} \rightarrow \text{Type}$, i.e., the Dijkstra state monad from the §1 with $\text{state} = \text{int}$.

Using the ST' computation type in F^* , one can specify for `incr` the type $\text{unit} \rightarrow ST' \text{unit } (\lambda \text{ post } s0 \rightarrow \forall s1. s1 > s0 \implies \text{post } ((), s1))$. That is, to prove any postcondition `post` of `incr`, it suffices to prove $\text{post } ((), s1)$ for any $s1$ greater than $s0$, the initial state—this is the statement of our goal. The proof in F^* currently involves:

- (1) As discussed already in §1, one must define $WP_ST t$, its return and bind combinators, proving that these specifications are sound with respect to the operational semantics of state.
- (2) The primitive effectful actions, `get` and `put` are assumed to have the types below—again, these types must be proven sound with respect to the operational semantics of F^* .

```
get : unit → ST' int (λ post s0 → post (s0, s0))
put : x:int → ST' unit (λ post _ → post ((), x))
```

- (3) Following the rule for sequential composition sketched in §1, F^* uses the specifications of `get` and `put` to compute $\text{bind_ST_WP } \text{wp.get } (\lambda x \rightarrow \text{wp.put } (x + 1))$ as the WP of `incr`, which reduces to $\lambda \text{ post } s0 \rightarrow \text{post } ((), s0 + 1)$.
- (4) The final step requires proving that the computed WP is at least as weak as the specified goal, which boils down to showing that $s0 + 1 > s0$, which F^* and Z3 handle automatically.

The first two steps above correspond to adding a new effect to F^* . The cost of this is amortized by the much more frequent and relatively automatic steps 3 and 4. However, adding a new effect to F^* is currently an expert activity, carried out mainly by the language designers themselves. This is in large part because the first two steps above are both tedious and highly technical: a dangerous mixture that can go wrong very easily.

Our primary goal is to simplify those first two steps, allowing effects to be added to F^* more easily and with fewer meta-level arguments to trust. Besides, although F^* supports customization of its effect system, it only allows programmers to specify refinements of a fixed set of existing effects inherited from ML, namely, state, exceptions, and divergence. For example, an F^* programmer can refine the state effect into three sub-effects for reading, writing, and allocation; but, she cannot add a new effect like alternative combinations of state and exceptions, non-determinism, continuations, etc. We aim for a more flexible, trustworthy mechanism for extending F^* beyond the primitive effects it currently supports. Furthermore, we wish to place reasoning about terminating effectful programs on

an equal footing with pure ones, supporting mixtures of intrinsic and extrinsic proofs for both.

2.3 CPS'ing monads to Dijkstra monads

Instead of manually specifying WP_ST , we program a traditional ST monad and derive WP_ST using a CPS transform. In §4.1 we formally present DM, a simply typed language in which to define monadic effects. DM itself contains a single primitive identity monad τ , which (as will be explained shortly) is used to control the CPS transform. We have implemented DM as a subset of F^* , and for the informal presentation here we use the concrete syntax of our implementation. What follows is an unsurprising definition of a state monad $\text{st } a$, the type of total functions from s to identity computations returning a pair $(a * s)$.

```
let st a = s → τ(a * s)
let return (x:a) : st a = λ s0 → x, s0
let bind (f:st a) (g:a → st b) : st b = λ s0 → let x,s1 = f s0 in g x s1
let get () : st a = λ s0 → s0, s0
let put (x:s) : st unit = λ _ → (), x
```

This being a subset of F^* , we can use it to prove that this definition is indeed a monad: proofs of the three monad laws for st are discharged automatically by F^* below (feq is extensional equality on functions, and `assert p` requests F^* to prove p statically). Other identities relating combinations of `get` and `put` can be proven similarly.

```
let right_unit_st (f:st α) = assert (feq (bind f return) f)
let left_unit_st (x:α) (f:(α → st β)) = assert (feq (bind (return x) f) (f x))
let assoc_st (f:st α) (g:(α → st β)) (h:(β → st γ))
  = assert (feq (bind f (λ x → bind (g x) h)) (bind (bind f g) h))
```

We then follow a two-step recipe to add an effect like st to F^* :

Step 1 To derive the Dijkstra monad variant of st , we apply a selective CPS transformation called the $\star$ -translation (§4); first, on type $\text{st } a$; then, on the various monadic operations. CPS'ing only those arrows that have τ -computation co-domains, we obtain:

```
(st a)*    = a → ((a * s) → Type) → Type
return*   = λ x s0 post → post (x, s0)
bind*     = λ f g s0 post → f s0 (λ(x,s1) → g x s1 post)
get*      = λ () s0 post → post (s0, s0)
put*      = λ x _ post → post ((), x)
```

Except for a reordering of arguments, the terms above are identical to the analogous definitions for WP_ST . We prove that the $\star$ -translation preserves equality: so, having shown the monad laws for $\text{st } a$, we automatically obtain the monad laws for $(\text{st } a)^*$. We also prove that every predicate transformer produced by the $\star$ -translation is monotone (it maps weaker postconditions to weaker preconditions) and conjunctive (they distribute over conjunctions and universals, i.e., infinite conjunctions, on the postcondition).

Step 2 The $\star$ -translation yields a predicate transformer semantics for a new monadic effect, however, we still need a way to extend F^* with the computational behavior of the new effect. For this, we define a second translation, which elaborates the definitions of the new monad and its associated actions to Pure computations in F^* . A first rough approximation of what we prove is that for a well-typed DM computation $e : \tau t$, its elaboration $\underline{e}$ has type $\text{Pure } \underline{t} \text{ e}^*$ in EMF^* .

The first-order cases are particularly simple: for example, `return` has type $x:a \rightarrow \text{Pure } a \text{ (return } x \text{)}$ in EMF^* ; and `get` = `get` has type $u:\text{unit} \rightarrow \text{Pure } s \text{ (get } u \text{)}$ in EMF^* . For a higher-order example, we sketch the elaboration of `bind` below, writing $\underline{\text{st } t} \text{ wp}$ for $s0:s \rightarrow \text{Pure } t \text{ (wp } s0)$:

```
bind : wpf:(st a)* → f:st a wpf
      → wpf:(a → (st b)*) → g:(x:a → st b wpf x)
      → st b (bind* wpf wpf)
      = λ wpf f wpf g s0 → let x, s1 = f s0 in g x s1
```


Intuitively, a function in DM (like `bind`) that abstracts over computations (`f` and `g`) is elaborated to a function (`bind`) in EMF^* that abstracts both over those computations (`f` and `g` again, but at their elaborated types) as well as the WP specifications of those computations (`wpf` and `wpg`). The result type of `bind` shows that it returns a computation whose specification matches `bind*`, i.e., the result of the CPS'ing $\star$ -translation.

In other words, the WPs computed by F^* for monads implemented as Pure programs corresponds exactly to what one gets by CPS'ing the monads. At first, this struck us as just a happy coincidence, although, of course, we now know that it must be so. We see our proof of this fact as providing a precise characterization of the close connection between WPs and CPS transformations.

2.4 Reify and reflect, for abstraction and proving

Unlike prior F^* formalizations which included primitive exception and state effects, the only primitive monad in EMF^* is for Pure computations.¹ Although the translations from DM yield pure definitions of monads in F^* , programming directly against those pure implementations is undesirable, since this may break abstractions. For instance, consider an integer-state monad whose state is expected to monotonically increase: revealing its representation as a pure term makes it hard to enforce this invariant. We rely on Filinski's (1994) monadic reflection for controlling abstraction.

Continuing our example, introducing the state effect in F^* produces a new computation type $\text{ST} (a:\text{Type}) (wp: (st\ a)^*)$ and two coercions

```
reify : ST a wp → s0:s → Pure (a * s) (wp s0)
reflect : (s0:s → Pure (a * s) (wp s0)) → ST a wp
```

The `reify` coercion reveals the representation of an ST computation as a Pure function, while `reflect` encapsulates a Pure function as a stateful computation. As we will see in subsequent sections, in some cases to preserve abstractions, one or both of these coercions will need to be removed, or restricted in various ways.

To introduce the actions from DM as effectful actions in F^* , we reflect the pure terms produced by the elaboration from DM to EMF^* , obtaining actions for the newly introduced computation type. For example, after reflection the actions `get` and `put` appear within F^* at the types below:

```
get : unit → ST s (get* ())
put : s1:s → ST unit (put* s1)
```

As in §2.2, we can still program stateful functions and prove them intrinsically, by providing detailed specifications to augment their definitions—of course, the first two steps of the process there are now automatic. However, we now have a means of doing extrinsic proofs by reifying stateful programs, as shown below (taking `s=int`).

```
let StNull a = ST a (λ s0 post → ∀x. post x)
let incr _ : StNull unit = let n = get() in put (n + 1)
let incr_increases (s0:s) = assert (snd (reify (incr()) s0) = s0 + 1)
```

The `StNull` unit annotation on the second line above gives a weak specification for `incr`. However, later, when a particular property of `incr` is required, we can recover it by reasoning extrinsically about the reification of `incr()` as a pure term.

¹ We leave divergence out of scope of the present work as a relatively orthogonal concept. We envisage adding divergence to EMF^* and DM as a second primitive effect in the future, with divergent computations interpreted in a partial correctness semantics with only intrinsic proving available. We do not foresee any significant difficulties in doing this, following the treatment of divergence of Swamy et al. (2016). We expect this to provide partial-correctness Dijkstra monads for free.

2.5 Combining monads: state and exceptions, in two ways

To add more effects to F^* , one can simply repeat the methodology outlined above. For instance, one can use DM to define $\text{exn}\ a = \text{unit} \rightarrow \tau(\text{option}\ a)$ in the obvious way (the `unit` is necessary, cf. §4.1), our automated two-step recipe extends F^* with an effect for terminating programs that may raise exceptions. Of course, we would like to combine the effects to equip stateful programs with exceptions and, here, we come to a familiar fork in the road.

State and exceptions can be combined in two mutually incompatible ways. In DM, we can define both $\text{stexn}\ a = s \rightarrow \tau((\text{option}\ a) * s)$ and $\text{exnst}\ a = s \rightarrow \tau(\text{option}\ (a * s))$. The former is more familiar to most programmers: raising an exception preserves the state; the latter discards the state when an exception is raised, which though less common, is also useful. We focus first on `exnst` and then discuss a variant of `stexn`.

Relating `st` and `exnst` Translating `st` (as before) and `exnst` to F^* gives us two unrelated effects `ST` and `ExnST`. To promote `ST` computations to `ExnST`, we define a lift relating `st` to `exnst`, their pure representations in DM, and prove that it is a monad morphism.

```
let lift (f:st a) : exnst a = λ s0 → Some (f s0)
let lift_is_an_st_exnst_morphism =
  assert (∀ x. feq (lift (ST.return x)) (ExnST.return x));
  assert (∀ f g. feq (lift (ST.bind f g)) (ExnST.bind (lift f) (λ x → lift (g x))))
```

Applying our two-step translation to `lift`, we obtain in F^* a computation-type coercion from $\text{ST}\ a\ wp$ to $\text{ExnST}\ a\ (\text{lift}^*\ wp)$. Through this coercion, and through F^* 's existing inference algorithm (Swamy et al. 2011, 2016), `ST` computations are implicitly promoted to `ExnST` computations whenever needed. In particular, the `ST` actions, `get` and `put`, are implicitly available with `ExnST`. All that remains is to define an additional action, `raise = λ() s0 → None`, which gets elaborated and reflected to F^* at the type $\text{unit} \rightarrow \text{ExnST}\ a\ (\lambda_p \rightarrow p\ \text{None})$.

`ExnST` programs in F^* can be verified intrinsically and extrinsically. For an intrinsic proof, we show `div_intrinsic` below, which raises an exception on a divide-by-zero. To prove it, we make use of an abbreviation `ExnSt a pre post`, which lets us write specifications using pre- and postconditions instead of predicate transformers, which can be more convenient—the F^* keywords, `requires` and `ensures` are only there for readability and have no semantic content.

```
let ExnSt a pre post = ExnST a (λ s0 p →
  pre s0 ∧ ∀x. post s0 x ⇒ p x)
let div_intrinsic i j : ExnSt int
  (requires (λ _ → True))
  (ensures (λ s0 x → match x with
    | None → j=0
    | Some (z, s1) → s0 = s1 ∧ j <> 0 ∧ z = i / j))
  = if j=0 then raise () else i / j
```

Alternatively, for an extrinsic proof, we give a weak specification for `div_extrinsic` and verify it by reasoning about its reified definition separately. This time, we add a call to `incr` in the `ST` effect in case of a division-by-zero. F^* 's type inference lifts `incr` to `ExnST` as required by the context. However, as the proof shows, the `incr` has no effect, since the `raise` that follows it discards the state.

```
let ExnStNull a = ExnST a (λ s0 post → ∀x. post x)
let div_extrinsic i j : ExnStNull int = if j=0 then (incr(); raise ()) else i / j
let lemma_div_extrinsic i j =
  assert (match reify (div_extrinsic i j) 0 with
    | None → j = 0
    | Some (z, 0) → j <> 0 ∧ z = i / j)
```

Using `reify` and `reflect` we can also build exception handlers, following ideas of Filinski (Filinski 1999). For example, in `try_div` below, we use a handler and (under-)specify that it never raises an exception.

```

let try_div i j : ExnSt int
  (requires (λ _ → True))
  (ensures (λ _ x → Option.isSome x))
= reflect (λ s0 → match reify (div_intrinsic i j) s0 with
  | None → Some (0, s0)
  | x → x)

```

More systematically, we can first program a Benton and Kennedy (2001) exception handler in DM, namely, as a term of type

$$\text{exnst } a \rightarrow (\text{unit} \rightarrow \text{exnst } b) \rightarrow (a \rightarrow \text{exnst } b) \rightarrow \text{exnst } b$$

and then translate it to F^* , thereby obtaining a weakest precondition rule for it for free. More generally, adapting Plotkin and Pretnar’s algebraic effect handlers (Plotkin and Pretnar 2009) to user-defined monads m , handlers can be programmed in DM as terms of type

$$m \ a \rightarrow (m \ b \rightarrow b) \rightarrow (a \rightarrow b) \rightarrow b$$

and then imported to F^* . We leave a more thorough investigation of such effect handlers for Dijkstra monads to the future.

An exception-counting state monad: stexnC For another combination of state and exceptions, we define stexnC , which in addition to combining state and exceptions (in the familiar way), also introduces an additional piece of integer state to count the number of exceptions that are raised. In DM, (omitting the standard return and bind) we write:

```

let stexnC a = (s * int) → τ(option a * (s * int))
let raise () = λ (s, n) → None, (s, n + 1)
let lift (f:st a) : stexnC a = λ (s, n) → let x, s1 = f s in Some x, (s1, n)

```

Notice that `raise` increments a counter. Adding StExnC to F^* proceeds as before. But, we need to be a bit careful with how we use reflection. In particular, an implicit invariant of stexnC is that its second state cell monotonically increases and actually counts the number of raised exceptions. If a programmer is allowed to reflect any $(s * \text{int}) \rightarrow \text{Pure}(\text{option } a * (s * \text{int}))$ wp into an StExnC computation, then this invariant can be broken. Programmers can rely on F^* ’s module system to simply forbid the use of StExnC.reflect in client modules. Depending on the situation, the module providing the effect may still reveal a restricted version of the reflect operator to a client, e.g., we may only provide `reflect_increasing` to clients, which only supports reflecting computations whose exception counter does not decrease. Of course, this only guarantees that the counter over-approximates the number of exceptions raised, which may or may not be acceptable.

```

let reflect_increasing (f: (s * int) → Pure (option a * (s * int))) wp
: StExnC a (λ (s0, n) post →
  wp s0 (λ (s1, n1) → post (s1, n1) ∧ n1 ≥ n0))
= reflect f

```

The standard combination of state and exceptions (i.e., stexnC) was already provided primitively in F^* . The other two combinations shown here were not previously supported, since F^* only allowed primitive effects. In the next two subsections, we present encodings of two other user-defined effects: a dynamic information-flow control monitor (§2.6) and continuations (§2.7).

2.6 Information-flow control

Information-flow control (Sabelfeld and Myers 2006) is a paradigm in which programs are deemed secure when one can prove that its behavior observable to an adversary is independent of the secrets the program may manipulate, i.e., they are *non-interferent*. Monadic reification allows us to prove non-interference properties directly, by relating multiple runs of an effectful program (Benton 2004). For example, take the simple stateful program below:

```

let ifc h = if h then (incr(); let y = get() in decr(); y) else get() + 1

```

It is easy to prove this program non-interferent via the extrinsic, relational proof below, which states that regardless of its secret input ($h0, h1$), ifc when run in the same public initial state ($s0$) produces identical public outputs. This generic extrinsic proof style is in contrast to Barthe et al. (2014), whose rF^* is a custom extension to F^* supporting only intrinsic relational proofs.

```

let ni_ifc = assert (∀ h0 h1 s0. reify (ifc h0) s0 = reify (ifc h1) s0)

```

Aside from such relational proofs, with user-defined effects, it is also possible to define monadic, dynamic information-flow control monitors in DM, deferring non-interference checks to runtime, and to reason about monitored programs in F^* . Here’s a simplified example, inspired by the floating label approach of LIO (Stefan et al. 2011). For simplicity, we take the underlying monad to be exnst , where the state is a security label from a two-point lattice that represents the secrecy of data that a computation may have observed so far.

```

type label = Low | High
let difc a = label → τ(option (a * label))

```

Once added to F^* , we can provide two primitive actions to interface with the outside world, where DIFC is the effect corresponding to `difc`. Importantly, writing to a public channel using `write Low` when the current label is `High` causes a dynamic failure signaling a potential Leak of secret information.

```

let join l1 l2 = match l1, l2 with | _, High | High, _ → High | _, Low
val read : l:label → DIFC bool (λ l0 p → ∀ b. p (Some (b, join l0 l)))
let flows l1 l2 = match l1, l2 with | High, Low → false | _ → true
val write : l:label → bool → DIFC unit (λ l0 p →
  if flows l0 l then p (Some ((), l0)) else p None)

```

As before, it is important to not allow untrusted client code to reflect on DIFC , since that may allow it to declassify arbitrary secrets. Arguing that DIFC soundly enforces a form of termination-insensitive non-interference requires a meta-level argument, much like that of Stefan et al. (2011).

We can now write programs like the one below, and rely on the dynamic checks to ensure they are secure.

```

let b1, b2 = read Low, read Low in write Low (b1 && b2)
let b3 = read High in write High (b1 || b3); write Low (xor b3 b3)

```

In this case, we can also prove that the program fails with a `None` at the last `write Low`. In contrast to the relational proof sketched earlier, dynamic information-flow control is conservative: even though the last `write` reveals no information on the low channel, the monitor raises an error.

2.7 CPS’ing the continuation monad

As a final example before our formal presentation, we ask the irresistible question of whether we can get a Dijkstra monad for free for the continuation monad itself—indeed, we can.

We start by defining the standard continuation monad, `cont`, in DM. Being a subset of F^* , we can prove that it is indeed a monad. The equality we need for this proof is an extensional equality at higher order—we use F^* ’s refinement types to define `kont`, a variant of `cont` augmented with an extensional equality principle, and (automatically) prove the monad laws for `kont`.

```

let cont a = (a → τ ans) → τ ans
let return x = λ k → k x
let bind f g k = f (λ x → g x k)
(* kont: continuations with an extensional equality principle *)
let kont a = f:(cont a) {∀ k1 k2. feq k1 k2 ⇒ f k1 = f k2}
(* kont is a monad *)
let r_unit (f:kont a) = assert (feq (bind f return) f);
let l_unit (x:a) (f:(a → kont b)) = assert (feq (bind (return x) f) (f x))
let assoc (f:kont a) (g:a → kont b) (h:b → kont c) =
  assert (feq (bind f (λ x → bind (g x) h)) (bind (bind f g) h))

```

Following our two-step recipe, we derive the Dijkstra variant of `cont`, but first we define some abbreviations to keep the notation manageable. The type `kwp a` is the type of a predicate transformer specifying a continuation `a` $\rightarrow \tau \text{ans}$; and `kans` is the type of a predicate transformer of the computation that yields the final answer.

$$\begin{aligned} \text{kwp } a &= a \rightarrow \text{kans} = (a \rightarrow \tau \text{ans})^* \\ \text{kans} &= (\text{ans} \rightarrow \text{Type}) \rightarrow \text{Type} = (\tau \text{ans})^* \end{aligned}$$

Using these abbreviations, we show the $\star$ -translation of `cont`, `return` and `bind`. Instead of being just a predicate transformer, `(cont a)*` is a predicate-transformer transformer.

$$\begin{aligned} (\text{cont } a)^* &= \text{kwp } a \rightarrow \text{kans} \\ \text{return}^* &= \lambda(x:a) (\text{wp_k:kwp } a) \rightarrow \text{wp_k } x \\ \text{bind}^* &= \lambda f g (\text{wp_k:kwp } b) \rightarrow f (\lambda(x:a) \rightarrow g \times \text{wp_k}) \end{aligned}$$

For step 2, we show the elaboration of `return` and `bind` to F^* , using the abbreviation `kt a wp` for the type of the elaborated term `k`, where the DM term `k` is a continuation of type `a` $\rightarrow \tau \text{ans}$ and `wp=k*`. As illustrated in §2.3, elaborating higher-order functions from DM to F^* introduces additional arguments corresponding to the predicate transformers of abstracted computations.

$$\begin{aligned} \text{kt } a \text{ wp} &= x:a \rightarrow \text{Pure ans} (\text{wp } x) \\ \text{return} &: x:a \rightarrow \text{wpk:kwp } a \rightarrow \text{kt } a \text{ wpk} \rightarrow \text{Pure ans} (\text{return}^* \times \text{wpk}) \\ &= \lambda x \text{ wpk } k \rightarrow k \times \\ \text{bind} &: \text{wpf:(cont } a)^* \\ &\rightarrow f:(\text{wpk:kwp } a \rightarrow \text{kt } a \text{ wpk} \rightarrow \text{Pure ans} (\text{wpf } \text{wpk})) \\ &\rightarrow \text{wpg:(a} \rightarrow (\text{cont } b)^*) \\ &\rightarrow g:(x:a \rightarrow \text{wpk:kwp } b \rightarrow \text{kt } b \text{ wpk} \rightarrow \text{Pure ans} (\text{wpg } \times \text{wpk})) \\ &\rightarrow \text{wpk:kwp } b \\ &\rightarrow k:\text{kt } b \text{ wpk} \\ &\rightarrow \text{Pure ans} (\text{bind}^* \text{ wpf } \text{wpg } \text{wpk}) \\ &= \lambda \text{wpf } f \text{ wpg } g \text{ wpk } k \rightarrow f (\lambda x \rightarrow \text{wpg } \times \text{wpk}) (\lambda x \rightarrow g \times \text{wpk } k) \end{aligned}$$

In the case of `return`, we have one additional argument for the predicate transformer of the continuation `k`—the type of the result shows how `return` relates to `return*`. The elaboration `bind` involves many such additional parameters, but the main point to take away is that its specification is given in terms of `bind*`, using the predicate transformers `wpf`, `wpg`, `wpk` in place of the `f`, `g`, `k` computations. In both cases, the definitions of `return` and `bind` match their pre-images in DM aside from abstracting over and passing around the additional WP arguments.

To better see the monadic structure in the types of `return` and `bind` we repeat these types, but this time writing `cont a wp` for the type `wpk:kwp a` $\rightarrow \text{kt } a \text{ wpk} \rightarrow \text{Pure ans} (\text{wp } \text{wpk})$:

$$\begin{aligned} \text{return} &: x:a \rightarrow \text{cont } a (\text{return}^* \times) \\ \text{bind} &: \text{wpf:(cont } a)^* \rightarrow f:\text{cont } a \text{ wpf} \\ &\rightarrow \text{wpg:(a} \rightarrow (\text{cont } b)^*) \rightarrow g:(x:a \rightarrow \text{cont } b (\text{wpg } x)) \\ &\rightarrow \text{cont } b (\text{bind}^* \text{ wpf } \text{wpg}) \end{aligned}$$

3. Explicitly monadic F^*

We begin our formal development by presenting EMF^* , an explicitly typed, monadic core calculus intended to serve as a model of F^* . As seen above, the F^* implementation includes an inference algorithm (Swamy et al. 2016) so that source programs may omit all explicit uses of the monadic `return`, `bind` and `lift` operators. We do not revisit that inference algorithm here and leave as future work a formal proof that after inference, F^* terms can be elaborated into EMF^* (along the lines of the elaboration of Swamy et al. (2011)).

3.1 Syntax

Figure 1 shows the EMF^* syntax. We highlight several key features.

Expressions, types, WPs, and formulae are all represented uniformly as terms; however, to evoke their different uses, we often write *e* for expressions, *t* for types, *wp* for WPs, and ϕ for logical

Terms

$$\begin{aligned} e, t, \text{wp}, \phi &::= x \mid T \mid x:t\{\phi\} \mid \lambda x:t. e \mid x:t \rightarrow c \mid e_1 e_2 \\ &\mid \text{case}_t(e \text{ as } y) x.e_1 x.e_2 \mid \text{run } e \mid \text{reify } e \\ &\mid \text{reflect } e \mid M.\text{lift}_{M'} t \text{ wp } e \mid F.\text{act } \bar{e} \\ &\mid M.\text{return } t \mid M.\text{bind } t_1 t_2 \text{ wp}_1 e_1 \text{ wp}_2 x.e_2 \end{aligned}$$

Computation types

$$c ::= \text{Tot } t \mid M t \text{ wp} \text{ where } M \in \{\text{Pure}, F\}$$

Signatures of monadic effects and lifts

$$\begin{aligned} S &::= D \mid S, D \mid S, L \\ D &::= F \left\{ \begin{array}{ll} \text{repr} &= t \quad ; \quad \text{wp_type} = t \\ \text{return} &= e \quad ; \quad \text{return}^* = \text{wp} \\ \text{bind} &= e \quad ; \quad \text{bind}^* = \text{wp} \\ \text{act}_j &= e \quad ; \quad \text{act}_j^* = \frac{x_j:t_j}{x_j:t_j} \rightarrow c_j \end{array} \right\} \\ L &::= \{ M.\text{lift}_{M'} = e; M.\text{lift}_{M'}^* = \text{wp} \} \end{aligned}$$

Figure 1. Syntax of EMF^*

formulae. Terms include variables (*x, y, a, b, w* etc.); refinement types $x:t\{\phi\}$; λ abstractions; dependent products with computation-type co-domains, $x:t \rightarrow c$ (with the sugar described in §2); and applications. Constants *T* include Type_i , the *i*th level from a countable hierarchy of predicative universes.² We also include constants for non-dependent pairs and disjoint unions; the former are eliminated using `fst` and `snd` (also constants), while the latter are eliminated using `caset` (*e as y*) *x.e₁* *x.e₂*, which is standard dependent pattern matching with an explicit return type *t* and a name for the scrutinee *y*, provided only when the dependency is necessary.

Computation types (*c*) include `Tot t`, the type of total *t*-returning terms, and `M t wp`, the type of a computation with effect *M*, return type *t*, and behavior specified by the predicate transformer *wp*. Let *M* range over the `Pure` effect as well as user-defined effects *F*.

Explicit monadic returns, binds, actions, lifts, reify, and reflect. `M.return` and `M.bind` are the monad operations for the effect *M*, with explicit arguments for the types and predicate transformers. `M.liftM'` *t wp e* lifts the *e* : *M t wp* to *M'*. A fully applied *F* action is written `F.act` $\bar{e}$. The `reify` and `reflect` operators are for monadic reflection, and `run` coerces a `Pure` computation to `Tot`.

Signatures for user-defined effects EMF^* is parameterized by a signature *S*. A user-defined effect *F t wp* is specified using *D*, the result of translating a DM monad. A definition *D* is a record containing several fields: *repr* is the type of an *F* computation reified as a pure term, *wp_type* is the type of the *wp* argument to *F*; *return*, *bind*, and *act_j* are EMF^* expressions, and *return^{*}*, *bind^{*}*, and *act_j^{*}* are EMF^* WPs (*act_j* is the *j*th action of *F*). We use *S.F.return* to denote the lookup of the *return* field from *F*'s definition in the signature *S*, and similar notation for the other fields.

² We have yet to model F^* 's universe polymorphism, making the universes in EMF^* less useful than the ones in F^* . Lacking universes polymorphism, we restrict computation to have results in Type_0 . A simple remediation would be replicate the monad definitions across the universe levels.

$\frac{\text{T-RETURN} \quad S; \Gamma \vdash e : \text{Tot } t}{S; \Gamma \vdash M.\text{return } t \ e : M \ t \ (S.M.\text{return}^* \ t \ e)}$	
$\text{T-BIND} \quad \frac{S; \Gamma \vdash t_2 : \text{Type}_0 \quad S; \Gamma \vdash wp_2 : x.t_1 \rightarrow S.M.\text{wp.type } t_2 \quad S; \Gamma \vdash e_1 : M \ t_1 \ wp_1 \quad S; \Gamma, x.t_1 \vdash e_2 : M \ t_2 \ (wp_2 \ x)}{S; \Gamma \vdash M.\text{bind } t_1 \ t_2 \ wp_1 \ e_1 \ wp_2 \ x.e_2 : M \ t_2 \ (S.M.\text{bind}^* \ t_1 \ t_2 \ wp_1 \ wp_2)}$	
$\text{T-LIFT} \quad \frac{S; \Gamma \vdash e : M \ t \ wp}{S; \Gamma \vdash M.\text{lift}_{M'} \ t \ wp \ e : M' \ t \ (S.M.\text{lift}_{M'}^* \ wp)}$	$\text{T-ACT} \quad \frac{S.F.\text{act}^* = \bar{x}.t \rightarrow c \quad \forall i. S; \Gamma \vdash e_i : t_i}{S; \Gamma \vdash F.\text{act } \bar{e} : c[\bar{e}/\bar{x}]}$
$\text{T-REIFY} \quad \frac{S; \Gamma \vdash e : F \ t \ wp}{S; \Gamma \vdash \text{reify } e : \text{Tot } (S.F.\text{repr } t \ wp)}$	$\text{T-REFLECT} \quad \frac{S; \Gamma \vdash e : \text{Tot } (S.F.\text{repr } t \ wp)}{S; \Gamma \vdash \text{reflect } e : F \ t \ wp}$
$\text{T-RUN} \quad \frac{S; \Gamma \vdash e : \text{Pure } t \ wp \quad S; \Gamma \models \exists p. wp \ p}{S; \Gamma \vdash \text{run } e : \text{Tot } t}$	$\text{T-SUB} \quad \frac{S; \Gamma \vdash e : c' \quad S; \Gamma \vdash c' <: c}{S; \Gamma \vdash e : c}$
$\text{T-REFINE} \quad \frac{S; \Gamma \vdash t : \text{Type}_i \quad S; \Gamma, x.t \vdash \phi : \text{Type}_j}{S; \Gamma \vdash x.t\{\phi\} : \text{Type}_i}$	$\text{C-PURE} \quad \frac{S; \Gamma \vdash t : \text{Type}_0 \quad S; \Gamma \vdash wp : (t \rightarrow \text{Type}_0) \rightarrow \text{Type}_0}{S; \Gamma \vdash \text{Pure } t \ wp : \text{Type}_0}$

Figure 2. Selected typing rules for EMF*

For example, for the ST monad from §2.3, we have³:

$ST\{$	$wp.\text{type} = \lambda a.s \rightarrow (a * s \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$
	$\text{repr} = \lambda a.w.s_0:s \rightarrow \text{Pure } (a * s) \ (w \ s_0)$
	$\text{return} = \lambda a.\text{return}$
	$\text{return}^* = \lambda a.\text{return}^*$
	$\text{bind} = \lambda a \ b.\text{bind}$
	$\text{bind}^* = \lambda a \ b.\text{bind}^*$
	$\text{get} = \text{get}$
	$\text{get}^* = \text{get}^*$
	$\text{put} = \text{put}$
	$\text{put}^* = \text{put}^* \}$

where (as described in §2.3) $\text{return} : a \rightarrow x.a \rightarrow \text{repr } a \ (\text{return}^* \ a \ x)$; and similarly for bind , get , and put .

In addition to the monad definitions D , the signature S contains the definitions of lifts that contain an EMF* expression and an EMF* WP. We use notations $S.M.\text{lift}_{M'}$ and $S.M.\text{lift}_{M'}^*$ to look these up in S . Finally, the signature always includes a fixed partial definition for the Pure monad, only containing the following definitions:

$\text{Pure}\{$	$wp.\text{type} = \lambda a:\text{Type}_0. (a \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$
	$\text{return}^* = \lambda a:\text{Type}_0. \lambda x.a. \lambda p:(a \rightarrow \text{Type}_0). p \ x$
	$\text{bind}^* = \lambda a. \lambda b. \lambda w_1. \lambda w_2. \lambda p. w_1 \ (\lambda x. (w_2 \ x) \ p) \}$

The other fields are not defined, since Pure is handled primitively in the EMF* dynamic semantics (§3.3).

3.2 Static semantics

The expression typing judgment in EMF* has the form $S; \Gamma \vdash e : c$, where Γ is the list of bindings $x : t$ as usual. Selected rules for the judgment are shown in Figure 2. In the rules, we sometimes write $S; \Gamma \vdash e : t$ as an abbreviation for $S; \Gamma \vdash e : \text{Tot } t$.

Monadic returns, binds, lifts, and actions. Rules T-RETURN, T-BIND, and T-LIFT simply use the corresponding wp specifica-

³ We use sans serif font for the actual field values.

$\text{S-TOT} \quad \frac{S; \Gamma \vdash t' <: t}{S; \Gamma \vdash \text{Tot } t' <: \text{Tot } t}$	$\text{S-PURE} \quad \frac{S; \Gamma \vdash t' <: t \quad S; \Gamma \models \forall p. wp \ p \Rightarrow wp' \ p}{S; \Gamma \vdash \text{Pure } t' \ wp' <: \text{Pure } t \ wp}$
$\text{S-F} \quad \frac{S; \Gamma \vdash S.F.\text{repr } t' \ wp' <: S.F.\text{repr } t \ wp}{S; \Gamma \vdash F \ t' \ wp' <: F \ t \ wp}$	$\text{S-PROD} \quad \frac{S; \Gamma \vdash t <: t' \quad S; \Gamma, x : t \vdash c' <: c}{S; \Gamma \vdash x:t \rightarrow c' <: x:t \rightarrow c}$
$\text{S-REFINEL} \quad \frac{S; \Gamma \vdash x:t\{\phi\} <: t}{S; \Gamma \vdash t <: x:t\{\phi\}}$	$\text{S-REFINER} \quad \frac{S; \Gamma, x : t \models \phi}{S; \Gamma \vdash t <: x:t\{\phi\}}$
	$\text{S-CONV} \quad \frac{S \vdash t' \rightarrow^* t \vee S \vdash t \rightarrow^* t'}{S; \Gamma \vdash t' <: t}$

Figure 3. Selected subtyping rules for EMF*

tion from the signature for M to compute the final wp . For example, in the case of the ST monad from §2.3, $S.ST.\text{return}^* \ t = \lambda x.t. \lambda s_0:s. \lambda \text{post.post} \ (x, s_0)$. Rule T-ACT is similar; it looks up the type of the action from the signature, and then behaves like the standard function application rule.

Monadic reflection and reification. Rules T-REIFY and T-REFLECT are dual, coercing between a computation type and its underlying pure representation. Rule T-RUN coerces e from type $\text{Pure } t \ wp$ to $\text{Tot } t$. However, since the Tot type is unconditionally total, the second premise of the rule checks that the wp is satisfiable.

Refinements, computations types, and proof irrelevance. EMF*'s refinement and computation types include a form of proof irrelevance. In T-REFINE, the universe of $x:t\{\phi\}$ is determined by the universe of t alone, since a witness for the proposition ϕ is never materialized. Refinement formulas ϕ and wps are manipulated using an entailment relation, $S; \Gamma \models \phi$, for a proof-irrelevant, classical logic where all the connectives are “squashed” (Nogin 2002), e.g., $p \wedge q$ and $p \Rightarrow q$ from §2, are encoded as $x:\text{unit}\{p * q\}$ and $x:\text{unit}\{p \rightarrow q\}$, and reside in Type_0 . Similar to T-REFINE, in C-PURE, the universe of a computation type is determined only by the result type. Since the wp is proof irrelevant, the use of Type_0 in the type of wp is quite natural, because its proof content is always squashed.

Subsumption and subtyping judgment. T-SUB is a subsumption rule for computations, which makes use of the two judgments $S; \Gamma \vdash c <: c'$ and $S; \Gamma \vdash t <: t'$, shown (selectively) in Figure 3. Rule S-PURE checks that $t' <: t$, and makes use of the $S; \Gamma \models \phi$ relation to check that wp is stronger than wp' , i.e. for all postconditions, the precondition computed by wp implies the precondition computed by wp' .

For user-defined monads F , the subtyping check delegates to their underlying representation $S.F.\text{repr}$. Rule S-PROD is the standard dependent function subtyping. Rule S-REFINEL permits dropping the refinement from the subtype, and rule S-REFINER allows subtyping to a refinement type, if we can prove the formula ϕ for an arbitrary x . Finally, rule S-CONV states that the beta-convertible types are subtypes of each other ($S \vdash t \rightarrow^* t'$ is the small-step evaluation judgment, introduced in the next section).

3.3 EMF* dynamic semantics

We now turn to the dynamic semantics of EMF*, which is formalized as a strong small-step reduction relation. Evaluation context are defined as follows:

$E ::=$	$\bullet \mid \lambda x.t.E \mid E \ e \mid e \ E \mid \text{run } E \mid \text{reify } E \mid \text{reflect } E$
	$\mid M.\text{bind } t_1 \ t_2 \ wp_1 \ E \ wp_2 \ x.e_2 \mid M.\text{return } t \ E$
	$\mid M.\text{lift}_{M'} \ t \ wp \ E \mid F.\text{act } \bar{e} \ E \ \bar{e}' \mid \text{case}_t(E \text{ as } _) \ x.e_1 \ x.e_2$
	$\mid \text{case}_t(e \text{ as } _) \ x.E_1 \ x.E_2 \mid \text{case}_t(e \text{ as } _) \ x.e_1 \ x.E_2$

$$\begin{array}{c}
\text{R-APP} \quad \frac{}{S \vdash (\lambda x:t.e) e' \longrightarrow e[e'/x]} \quad \text{R-RUN} \quad \frac{}{S \vdash \text{run } (\text{Pure.return } t \ e) \longrightarrow e} \\
\\
\text{R-PUREBIND} \quad \frac{}{S \vdash \text{Pure.bind } t_1 \ t_2 \ wp_1 \ (\text{Pure.return } t \ e_1) \ wp_2 \ x.e_2 \longrightarrow e_2[e_1/x]} \\
\\
\text{R-REIFYRET} \quad \frac{}{S \vdash \text{reify } (F.\text{return } t \ e) \longrightarrow S.F.\text{return } t \ e} \quad \text{R-REIFYREFLECT} \quad \frac{}{S \vdash \text{reify } (\text{reflect } e) \longrightarrow e} \\
\\
\text{R-REIFYBIND} \quad \frac{e' = S.F.\text{bind } t_1 \ t_2 \ wp_1 \ (\text{reify } e_1) \ wp_2 \ x.(\text{reify } e_2)}{S \vdash \text{reify } (F.\text{bind } t_1 \ t_2 \ wp_1 \ e_1 \ wp_2 \ x.e_2) \longrightarrow e'} \\
\\
\text{R-REIFYACT} \quad \frac{}{S \vdash \text{reify}(F.\text{act } \bar{e}) \longrightarrow S.F.\text{act}[\bar{e}/\bar{x}]} \\
\\
\text{R-REIFYLIFT} \quad \frac{}{S \vdash \text{reify}(M.\text{lift}_{M'} \ t \ wp \ e) \longrightarrow S.M.\text{lift}_{M'} \ t \ wp \ (\text{reify } e)}
\end{array}$$

Figure 4. Dynamic semantics of EMF* (selected reduction rules)

The judgment has the form $S \vdash e \longrightarrow e'$. We show some selected rules in Figure 4. The main ideas of the judgment are: (a) the Tot terms reduce primitively in a standard manner, (b) `Pure.bind` is also given a primitive semantics, however (c) to β -reduce other monadic operations (binds, returns, actions, and lifts), they need to be reified first, which then makes progress using their underlying implementation in the signature.

Semantics for Pure terms. Rule R-PUREBIND reduces similarly to the usual β -reduction. For run e , the semantics first evaluates e to `Pure.return` $t \ e'$, and then run removes the `Pure.return` and steps to the underlying total computation e' via R-RUN.

Semantics for monadic returns and binds. Rule R-REIFYBIND looks up the underlying implementation $S.F.\text{bind}$ in the signature, and applies it to e_1 and e_2 but after reifying them so that their effects are handled properly. In a similar manner, rule R-REIFYRET looks up the underlying implementation $S.F.\text{return}$ and applies it to e . Note that in this case, we don't need to reify e (as we did in bind), because e is already a Tot term.

Semantics for monadic lifts and actions. Rules R-REIFYACT and R-REIFYLIFT also lookup the underlying implementations of the lifts and actions in the signature and use them. Rule R-REIFYLIFT in addition reifies the computation e . For lifts, the arguments $\bar{e}$ are already Tot.

3.4 EMF* metatheory

We prove several metatheoretical results for EMF*. First, we prove strong normalization for EMF* via a translation to the calculus of inductive constructions (CiC) (Paulin-Mohring 2015).

Theorem 1 (Strong normalization). *If $S; \Gamma \vdash e : c$ and CiC is strongly normalizing, then e is strongly normalizing.*

Proof. (sketch) The proof proceeds by defining a translation from EMF* to CiC, erasing refinements and WPs, inlining the pure implementations of each monad, and removing the reify and reflect operators. We show that this translation is a type-preserving, forward simulation. If CiC is strongly normalizing, then EMF* must also be, since otherwise an infinite reduction sequence in EMF* could not be matched by CiC, contradicting the forward simulation. $\square$

Theorem 2 (Subject Reduction). *If $S; \Gamma \vdash e : c$ and $S \vdash e \longrightarrow e'$, then $S; \Gamma \vdash e' : c$.*

This allows us to derive a total correctness property for the Pure monad saying that run-ing a Pure computation produces a value which satisfies all the postconditions that are consistent with the wp of the Pure computation.

Corollary 3 (Total Correctness of Pure). *If $S; \cdot \vdash e : \text{Pure } t \ wp$, then $\forall p. S; \cdot \vdash p : t \rightarrow \text{Type}_0$ and $S; \cdot \models wp \ p$, we have $S \vdash \text{run } e \longrightarrow^* v$ such that $S; \cdot \models p \ v$.*

For the user-defined monads F , we can derive their total correctness property by appealing to the total correctness of the Pure monad. For instance, for the ST monad from §2.3, we can derive the following corollary simply by using the typing of reify and Corollary 3.

Corollary 4 (Total Correctness of ST). *If $S; \cdot \vdash e : ST \ t \ wp$, then $\forall p, s_0. S; \cdot \vdash s_0 : s$, $S; \cdot \vdash p : t \times s \rightarrow \text{Type}_0$ and $S; \cdot \models wp \ s_0 \ p$, then $S \vdash \text{run } ((\text{reify } e) \ s_0) \longrightarrow^* v$ such that $S; \cdot \models p \ v$.*

3.5 Implementation in F*

The implementation of F* was relatively easy to adapt to EMF*. In fact, EMF* and DM and the translation between them were designed to match F*'s existing type system, as much as possible. We describe the main changes that were made.

User-defined non-primitive effects are, of course, the main new feature. Effect configurations closely match the D form from Figure 1, the main delta being that non-primitive effects include pure implementations or $M.\text{bind}$, $M.\text{return}$, $M.\text{lift}_{M'}$, etc.

Handling reify and reflect in the type-checker involved implementing the two relatively simple rules for them in Figure 2. A more significant change was made to F*'s normalization machinery, extending it to support rules that trigger evaluation for reified, effectful programs. In contrast, before our changes, F* would never reduce effectful terms. The change to the normalizer is exploited by F*'s encoding of proof obligations to an SMT solver—it now encodes the semantics of effectful terms to the solver, after using the normalizer to partially evaluate a reified effectful term to its pure form.

Running programs with user-defined effects is achieved by extracting it to OCaml, as is usual for F*, except we now inline the definitions of the underlying pure terms. Effects marked as primitive are extracted as before while making use of primitive effects in OCaml—this is formally justified in §5.

4. Dijkstra monads for free

This section formally presents DM, a language for defining effects by giving monads with their actions and lifts between them. Via a pair of translations, we export such definitions to EMF* as effect configurations. The first translation of a term e , a CPS, written e^* produces a predicate-transformer from DM term; the second one is an *elaboration*, $\bar{e}$, which produces an EMF* implementation of a DM term. The main result shows that for any DM term the result of the $\star$ -translation is in a suitable logical relation to the elaboration of the term, and thus a valid specification for this elaboration. We also show that the $\star$ -translation always produces monotonic and conjunctive predicates, properties that should always hold for WPs. Finally, we show that the $\star$ -translation preserves all equalities in DM, and thus translates DM monads into EMF* Dijkstra monads.

4.1 Source: DM effect definition language

The source language DM is a simply-typed lambda calculus augmented with an abstract monad τ , as in §2.3. The language is essentially that of Filinski (1994) with certain restrictions on allowed types to ensure the correctness of elaboration.

There are two effect symbols: n (non-effectful) and τ . The typing judgment is split accordingly, and ε ranges over both of them. Every

monadic term needs to be bound via **bind**_τ to be used.⁴ Functions can only take non-effectful terms as arguments, but may return a monadic result.

The set of DM types is divided into *A types*, *H types*, and *C types*, ranged over by *A*, *C*, and *H*, respectively. They are given by the grammar:

$$\begin{aligned} A &::= X \mid b \mid A \xrightarrow{n} A \mid A + A \mid A \times A \\ H &::= A \mid C \\ C &::= H \xrightarrow{\tau} A \mid H \xrightarrow{n} C \mid C \times C \end{aligned}$$

Here *X* ranges over type variables (needed to define monads) and *b* are base types. The τ -arrows represent functions with a monadic result, and our translations will provide WPs for these arrows. *A* types are referred to as “ τ -free”, since they contain no monadic operations. *C* types are inherently computational in the sense that they cannot be eliminated into an *A* type: every possible elimination will lead to a monadic term. They are referred to as “computational types”. *H* types are the union of both, and are called “hypothesis” types, as they represent the types of possible functional arguments. As an example, the state monad is represented as the type $S \xrightarrow{\tau} (X \times S)$, where *X* is a type variable and *S* is some type representing the state.

DM types do not include “mixed” $A \times C$ pairs, computational sums $C + H$, functions of type $C \xrightarrow{n} A$, or types with right-nested τ -arrows. We do allow nesting τ -arrows to the left, providing the generality needed for the continuation monad, and others. These restrictions are crafted to carefully match EMF*. Without them, our translations, would generate ill-typed or logically unrelated EMF* terms, and they do not appear to be severe in practice, as evidenced by the examples in §2.

The syntax for terms is (κ standing for constants):

$$\begin{aligned} e &::= x \mid e \mid \lambda x.H. e \mid \kappa(e, \dots, e) \\ &\quad \mid (e, e) \mid \mathbf{fst}(e) \mid \mathbf{snd}(e) \\ &\quad \mid \mathbf{inl}(e) \mid \mathbf{inr}(e) \mid \mathbf{case} \ e \ \mathbf{inl} \ x.A. \ e; \ \mathbf{inr} \ y.A. \ e \\ &\quad \mid \mathbf{return}_\tau e \mid \mathbf{bind}_\tau e \ \mathbf{to} \ x \ \mathbf{in} \ e \end{aligned}$$

Typing judgments have the forms $\Delta \mid \Gamma \vdash e : H ! n$ and $\Delta \mid \Gamma \vdash e : A ! \tau$, where Δ is a finite sequence of type variables and Γ is a normal typing context, whose types only use type variables from Δ . Here are some example rules:

$$\begin{array}{c} \frac{\Delta \mid \Gamma, x.H \vdash e : H' ! \varepsilon}{\Delta \mid \Gamma \vdash \lambda x.H. e : H \xrightarrow{\varepsilon} H' ! n} \quad \frac{\Delta \mid \Gamma \vdash f : H \xrightarrow{\varepsilon} H' ! n \quad \Delta \mid \Gamma \vdash e : H ! n}{\Delta \mid \Gamma \vdash f e : H' ! \varepsilon} \\ \frac{\Delta \mid \Gamma \vdash e : A ! n}{\Delta \mid \Gamma \vdash \mathbf{return}_\tau e : A ! \tau} \quad \frac{\Delta \mid \Gamma \vdash e_1 : A' ! \tau \quad \Delta \mid \Gamma, x.A \vdash e_2 : A' ! \tau}{\Delta \mid \Gamma \vdash \mathbf{bind}_\tau e_1 \ \mathbf{to} \ x \ \mathbf{in} \ e_2 : A' ! \tau} \end{array}$$

In these rules we implicitly assume that all appearing types are well-formed with respect to the grammar, e.g., one cannot form a function of type $C \xrightarrow{n} A$ by the abstraction rule.

As an example, $\mathbf{return}_\tau = \lambda x.X. \lambda s.S. \mathbf{return}_\tau(x, s)$ has type $X \xrightarrow{n} S \xrightarrow{\tau} (X \times S)$, using these rules.

When defining effects and actions, one deals (at a top level) with non-effectful *C* types ($C ! n$). We present our main results (Theorem 5 and Theorem 6) for terms typed in an empty Γ as these are the most interesting cases; of course to show them, we prove more general results.

4.2 The $\star$ -translation

The essence of the $\star$ -translation is to translate $\mathbf{return}_\tau e$ and $\mathbf{bind}_\tau e_1 \ \mathbf{to} \ x \ \mathbf{in} \ e_2$ to the returns and binds of the continuation

⁴ In this formalization, **bind** and **return** appear explicitly in source programs. When using our implementation, however, the user need not call **bind** and **return**; rather, they write programs in a direct style, and **let**-bindings are turned into **binds** as needed. §4.6 provides some details on the interpretation and elaboration of concrete F^* terms as DM terms.

monad. We begin by defining a translation H^* , that translates any *H* type to the type of its predicates by CPS’ing the τ -arrows. First, for any τ -free type *A*, A^* is essentially the identity, except we replace every arrow $\xrightarrow{n}$ by $\rightarrow$. Then, for computation types, we define:

$$\begin{aligned} (H \xrightarrow{n} C)^* &= H^* \rightarrow C^* \\ (C \times C')^* &= C^* \times C'^* \\ (H \xrightarrow{\tau} A)^* &= H^* \rightarrow (A^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0 \end{aligned}$$

Note that all arrows on the right hand side of the translation have a Tot codomain, as per our notational convention.

In essence, the codomains of τ -arrows are CPS’d into a WP, which takes as argument a predicate on the result and produces a predicate representing the “precondition”. All other constructs are just translated recursively: the real work is for the τ -arrows.

For example, for the state monad $S \xrightarrow{\tau} (X \times S)$, the $\star$ -translation produces the EMF* type $S \rightarrow (X \times S \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$. It is the type of predicates that map an initial state and a postcondition (on both result and state) into a proposition. Modulo isomorphism (of the order of the arguments and Curryng)⁵ this is exactly the type of WPs in current F^* ’s state monad (cf. §1, §2.3).

The two main cases for the $\star$ -translation for well-typed DM terms are shown below; every other case is simply a homomorphic application of $\star$ on the sub-terms.

$$\begin{aligned} (\mathbf{return}_\tau e)^* &= \lambda p.(A^* \rightarrow \text{Type}_0). p \ e^* \text{ when } \Delta \mid \Gamma \vdash e : A ! n \\ (\mathbf{bind}_\tau e_1 \ \mathbf{to} \ x \ \mathbf{in} \ e_2)^* &= \lambda p.(A^* \rightarrow \text{Type}_0). e_1^* (\lambda x.A. e_2^* p) \\ &\quad \text{when } \Delta \mid \Gamma, x.A \vdash e_2 : A' ! \tau \end{aligned}$$

Formally, the $\star$ -translation and elaboration are defined over a typing derivation, as one needs more information than what is present on the term. The $\star$ -translation of terms and types are related in the following sense, where we define the context Δ as $X_1 : \text{Type}_0, \dots, X_n : \text{Type}_0$ when $\Delta = X_1, \dots, X_n$ (we assume that type variables are also EMF* variables).

After translation, one can abstract over the variables in Δ to introduce the needed polymorphism in EMF*. This will also be the case for elaboration.

For example, for the previous definition of $\mathbf{return}_\tau$ we get the translation $\lambda x.X. \lambda s.S. \lambda p.(X \times S \rightarrow \text{Type}_0). p(x, s)$, which has the required transformer type: $S \rightarrow (X \times S \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$. It is what one would expect: to prove a postcondition *p* about $\mathbf{return}_\tau x$, one needs to prove $p(x, s)$ where *s* is the initial state.

Theorem 5 (well-typing of $\star$ -translation).

$$\Delta \mid \cdot \vdash e : C ! n \text{ implies } \Delta \vdash e^* : C^*.$$

4.3 Elaboration

Elaboration is merely a massaging of the source term to make it properly typed in EMF*. During elaboration, monadic operations are translated to those of the identity monad in EMF*, namely *Pure*.

Elaboration of types We define two elaboration translations for DM types, which produce the EMF* types of the elaborated expression-level terms. The first translation $\underline{A}$ maps an *A* type to a simple EMF* type, while the second one F_C maps a *C* type and a term *wp* of type C^* (a specification) into an EMF* computational product or pair. The $\underline{A}$ translation is the same as the CPS one, i.e., $\underline{A} = A^*$.

The F_C *wp* (where *wp* : C^*) translation is defined by:

$$\begin{aligned} (1) \quad F_{C \times C'} \ wp &=_{\text{def}} F_C (\mathbf{fst} \ wp) \times F_{C'} (\mathbf{snd} \ wp) \\ (2) \quad F_{C \xrightarrow{\varepsilon} H} \ wp &=_{\text{def}} w' : C^* \rightarrow F_C w' \rightarrow G_H^\varepsilon (wp \ w') \\ (3) \quad F_{A \xrightarrow{\tau} H} \ wp &=_{\text{def}} x : A \rightarrow G_H^\tau (wp \ x) \end{aligned}$$

⁵ One can tweak our translation to generate WPs that have the usual postcondition to precondition shape. However we found the current shape to be generally easier to work with.

(1) $\underline{x}$	$= x$	(5) $\underline{\text{fst}}(e)$	$= \text{fst } e$
(2) $\underline{\kappa}(e_1, \dots, e_n)$	$= \kappa \ e_1 \dots e_n$	(6) $\underline{\text{snd}}(e)$	$= \text{snd } e$
(3) $\underline{\lambda x:A. e}$	$= \lambda x:A. \underline{e}$	(7) $\underline{\text{inl}}(e)$	$= \text{inl } e$
(4) $\underline{\lambda x:C. e}$	$= \lambda x^w:C^*. \lambda x:F_C x^w. \underline{e}$	(8) $\underline{\text{inr}}(e)$	$= \text{inr } e$
(9) $\underline{e_1 e_2}$	$= \underline{e_1} \ \underline{e_2}$		$(\Delta \mid \Gamma \vdash e_2 : A!n)$
(10) $\underline{e_1 e_2}$	$= \underline{e_1} \ (e_2^* \text{ s}_\Gamma) \ \underline{e_2}$		$(\Delta \mid \Gamma \vdash e_2 : C!n)$
(11) $\underline{(e_1, e_2)}$	$= (\underline{e_1}, \underline{e_2})$		
(12) $\underline{\text{case } e \text{ inl } x:A_1. e_1; \text{ inr } y:A_2. e_2}$	$= \text{case}(\underline{e}) \ x. \underline{e_1} \ y. \underline{e_2}$		$(\Delta \mid \Gamma, x:A_1 \vdash e_1 : A! \varepsilon)$
(13) $\underline{\text{case } e \text{ inl } x:A_1. e_1; \text{ inr } y:A_2. e_2}$	$= \text{case}_{F_C} \text{ case}(z) \ x. (e_1^* \text{ s}_\Gamma) \ y. (e_2^* \text{ s}_\Gamma) (\underline{e} \text{ as } z) \ x. \underline{e_1} \ y. \underline{e_2}$		$(\Delta \mid \Gamma, x:A_1 \vdash e_1 : C!n)$
(14) $\underline{\text{return}_\tau e}$	$= \text{Pure.return } \underline{A} \ \underline{e}$		$(\Delta \mid \Gamma \vdash e : A! \tau)$
(15) $\underline{\text{bind}_\tau e_1 \text{ to } x:A \text{ in } e_2}$	$= \text{Pure.bind } \underline{A} \ \underline{A'} \ (e_1^* \text{ s}_\Gamma) \ \underline{e_1} \ (\lambda x:A^*. e_2^* \text{ s}_\Gamma) \ x. \underline{e_2}$		$(\Delta \mid \Gamma, x:A \vdash e_2 : A'! \tau)$

Figure 5. The elaboration of DM terms to EMF*

Here we use the notation that $G_C^n(wp) = F_C \ wp$ and $G_A^\tau(wp) = \text{Pure } \underline{A} \ wp$.

The main idea is that if an EMF* term e has type $F_C \ wp$, then wp is a proper specification of the final result. Putting pairs aside for a moment, this means that if one applies enough arguments e_i to e in order to eliminate it into a Pure computation, then $e \ \bar{e}_i : \text{Pure } A \ (wp \ \bar{s}_i)$, where each s_i is the specification for each e_i . This is naturally extended to pairs by the definition, defining the specification for a pair as a pair of proper specifications, as shown by case (1) above.

In case (2), the $w' : C^*$ arguments introduced by F are relevant for the higher-order cases, and serve the following purpose, as illustrated in §2.3 (for the translation of bind for the ST monad, and in §2.7 for the continuation monad): when taking computations as arguments, we first require their specification in order to be able to reason about them at the type level. Taking these specification arguments is also the only way for being WP-polymorphic in EMF*. Note that, according to the dependencies, the C^* argument is only used in the specifications, while we shall see in elaboration that only the $F_C \ wp$ is used in the expressions. When elaborating terms, we pass this specification as an extra argument where needed.

In case (3), when elaborating functions taking an argument of A type there is no need to take a specification, since the argument is completely non-effectful and can be used at both expression- and type-levels. Informally, a non-effectful term is its own specification.

Returning to our state monad example, the result of $F_{S \rightarrow (X \times S)} \ wp$ is $s:S \rightarrow \text{Pure } (X \times S) \ (wp \ s)$, i.e., the type of a function f which for any p , if one can prove $wp \ s \ p$, then $f \ s$ satisfies p .

Elaboration of terms is defined in Figure 5 and is, as expected, mostly determined by the translation of types. The translation is formally defined over typing derivations, however, for brevity, we present each translation rule simply on the terms, with the important side-conditions we rely on from the derivation shown in parenthesis. We describe only the most interesting cases.

Computational abstractions and applications (cases 4 and 10)

Case (4) translates a function with a computational argument $x:C$ to a function that expects two arguments, a WP $x^w:C^*$ and x itself, related to x^w at a suitably translated type. We track the association between x and x^w using a substitution s_Γ , which maps every computational hypothesis $x : C$ in Γ to x^w (of type C^*) in $\underline{\Gamma}$. In case (10), when passing a computation argument e_2 , we need to eliminate the double abstraction introduced in case (4), passing both $e_2^* \text{ s}_\Gamma$, i.e. the WP of e_2 (substituting free computation variables), and $\underline{e_2}$ itself.

Return and bind (cases 14 and 15) The last two rules show the translation of return and bind for τ to return and bind for Pure in EMF*. This is one of the key points: in the elaboration, we interpret the τ as the identity monad in EMF*, whereas in the $\star$ -translation, we interpret τ as the continuation monad. Theorem 6, our main

theorem, shows that EMF*'s WP computation in the Pure monad for $\underline{e}$ produces a WP that is logically related to the $\star$ -translation of e , i.e., WPs and the CPS coincide formally, at arbitrary order.

Theorem 6 (Logical relations lemma).

$\Delta \mid \cdot \vdash e : C!n$ implies $\underline{\Delta} \vdash \underline{e} : F_C \ e^*$

4.4 Monotonicity and conjunctivity

A key property of WPs is monotonicity: weaker postconditions should map to weaker preconditions. This is also an important F* invariant that allows for logical optimizations of WPs. Similarly, WPs are conjunctive: they distribute over conjunction and universal quantification in the postcondition. We show that any EMF* term obtained from the $\star$ -translation is monotonic and conjunctive, for higher-order generalizations of the usual definitions of these properties (Dijkstra 1997).

We first define a logical relation between EMF* terms $t_1 \lesssim_t t_2$, read “ t_1 stronger than t_2 at type t ” and producing an EMF* formula in Type_0 , by recursion on the structure of t :

$$\begin{aligned}
x \lesssim_{\text{Type}_0} y &= \text{def } x \Rightarrow y \\
x \lesssim_b y &= \text{def } x = y \\
x \lesssim_X y &= \text{def } x = y \\
f \lesssim_{t_1 \rightarrow t_2} g &= \text{def } \forall x_1, x_2 : t_1. x_1 \lesssim_{t_1} x_2 \Rightarrow f \ x_1 \lesssim_{t_2} g \ x_2 \\
x \lesssim_{t_1 \times t_2} y &= \text{def } \text{fst } x \lesssim_{t_1} \text{fst } y \wedge \text{snd } x \lesssim_{t_2} \text{snd } y \\
x \lesssim_{t_1 + t_2} y &= \text{def } (\exists v_1, v_2 : t_1, x = \text{inl } v_1 \wedge y = \text{inl } v_2 \wedge v_1 \lesssim_{t_1} v_2) \vee \\
&\quad (\exists v_1, v_2 : t_2, x = \text{inr } v_1 \wedge y = \text{inr } v_2 \wedge v_1 \lesssim_{t_2} v_2)
\end{aligned}$$

where b represents any EMF* base type and X any variable⁶.

For any type t that doesn't mention Type_0 , the relation reduces to extensional equality. The relation is only defined for the subset of EMF* types that are all-Tot and non-dependent. All types resulting from the $\star$ -translation will be in this subset, so this not a limitation for our purposes.

The $\lesssim$ relation is not reflexive. We say that a closed EMF* expression e for which $\cdot \vdash e : \text{Tot } t$ is *monotonic* when $\cdot \vdash e \lesssim_t e$. For the first-order WPs this coincides with the standard definition, and for higher-order predicates it gives a reasonable extension. For a first-order example, let's take the type of WPs for ST programs: $S \rightarrow (S \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$:

$$\begin{aligned}
&f \lesssim_{S \rightarrow (S \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} f \\
\equiv &\forall s_1, s_2. s_1 = s_2 \Rightarrow f \ s_1 \lesssim_{(S \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} f \ s_2 \\
\iff &\forall s. f \ s \lesssim_{(S \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} f \ s \\
\equiv &\forall s, p_1, p_2. p_1 \lesssim_{S \rightarrow \text{Type}_0} p_2 \Rightarrow f \ s \ p_1 \lesssim_{\text{Type}_0} f \ s \ p_2 \\
\iff &\forall s, p_1, p_2. (\forall s', p_1 \ s' \Rightarrow p_2 \ s') \Rightarrow (f \ s \ p_1 \Rightarrow f \ s \ p_2)
\end{aligned}$$

⁶ We can get a stronger result if we don't restrict the relation on type variables to equality and treat it abstractly instead. For our purposes this is not needed as we plan to instantiate type variables with predicate-free types

This is exactly the usual notion of monotonicity for imperative programs (Dijkstra 1997): “if p_2 is weaker than p_1 , then $f \text{ s } p_2$ is weaker than $f \text{ s } p_1$ for any s ”.

Since the $\lesssim$ relation is trivially preserved by application, monotonicity is too. Also, first-order predicates on base types (such as first-order pre-/postconditions) are always monotonic according to this definition.

We proved that the $\star$ -translation of any well-typed source term $e : C!n$ gives a monotonic e^* at the type C^* . This result is more general than it appears at a first glance: not only does it mean that WPs obtained by *any* defined return and bind are monotonic, but also for any action or function. Also, lifts between monads and other higher-level computations will preserve this monotonicity. Furthermore, the relation $\models$ in the conclusion of the theorem is EMF^* ’s validity judgment, i.e., we prove that these properties are actually provable within F^* without needing to rely on some meta-level reasoning.

Theorem 7 (Monotonicity of $\star$ -translation).

For any e and $C, \Delta \mid \cdot \vdash e : C!n$ implies $\Delta \models e^* \leq_{C^*} e^*$.

We give a similar higher-order definition of conjunctivity, and prove similar results ensuring the $\star$ -translation provides conjunctivity. The definition for conjunctivity is given by the following, where a describes the predicate-free types (including variables).

$$\begin{aligned} \mathbb{C}_{(a \rightarrow \text{Type}_0) \rightarrow \text{Type}_0}(w) &= \text{def } \forall p_1, p_2. w \text{ p}_1 \wedge w \text{ p}_2 = w (\lambda x. p_1 \text{ x} \wedge p_2 \text{ x}) \\ \mathbb{C}_a(x) &= \text{def } \text{true} \\ \mathbb{C}_{t_1 \rightarrow t_2}(f) &= \text{def } \forall x : t_1. \mathbb{C}_{t_1}(x) \Rightarrow \mathbb{C}_{t_2}(fx) \\ \mathbb{C}_{t_1 \times t_2}(p) &= \text{def } \mathbb{C}_{t_1}(\text{fst } p) \wedge \mathbb{C}_{t_2}(\text{snd } p) \end{aligned}$$

Again, the relation is not defined on all types, but it does include the image of the type-level $\star$ -translation, so it is enough for our purposes. This relation is also trivially preserved by application. We then prove:

Theorem 8 (Conjunctivity of $\star$ -translation).

For any e and $C, \Delta \mid \cdot \vdash e : C!n$ implies $\Delta \models \mathbb{C}_{C^*}(e^*)$

4.5 The $\star$ -translation preserves equality and monad laws

We define an equality judgment on source terms that is basically $\beta\eta$ -equivalence, augmented with the monad laws for the abstract τ monad. We show that the $\star$ -translation preserves this equality.

Theorem 9 (Preservation of equality by CPS).

If $\Delta \mid \cdot \vdash e_1 = e_2 : H!E$ then $\Delta \models e_1^* = e_2^*$.

Since the monad laws are equalities themselves, any source monad will be translated to a specification-level monad of WPs. This also applies to lifts: source monad morphisms are mapped to monad morphisms between Dijkstra monads.

4.6 Implementing the translations in F^*

We devised a prototype implementation of the two translations in F^* . Users define their monadic effects as F^* terms in direct style, as done in §2, and these definitions get automatically rewritten into DM. As explained in §2, instead of τ -arrows ($H \xrightarrow{\tau} A$), we use a distinguished F^* effect τ to indicate where the CPS should occur. The effect τ is defined to be an alias for F^* ’s Tot effect, which allows the programmer to reason extrinsically about the definitions and prove that they satisfy various properties within F^* , e.g., the monad laws. Once the definitions have been type-checked in F^* , another minimalistic type-checker kicks in, which has a twofold role. First, it ensures that the definitions indeed belong to DM, e.g., distinguishing A types from C types. Second, it performs bidirectional inference to distinguish monadic computations from pure computations, starting from top-level annotations, and uses this type information to automatically introduce **return _{τ}** and **bind _{τ}** as needed. For instance, in the st example from §2.3, the type-checker rewrites $x, s0$ into **return _{τ}** $(x, s0)$; and **let** $x, s1 = f \text{ s } 0$ in ...

into **bind _{τ}** $f \text{ s } 0$ **to** $x, s1$ **in** ...; and $g \text{ s } s1$ into **return _{τ}** $(g \text{ s } s1)$. The elaboration maps let-bindings in DM to let-bindings in F^* ; the general inference mechanism in F^* takes care of synthesizing the WPs, meaning that the elaboration, really, is only concerned about extra arguments for abstractions and applications.

Once the effect definition is rewritten to DM, our tool uses the $\star$ -translation and elaboration to generate the WP transformers for the Dijkstra monad, which previously would be written by hand. Moreover, using the $\lesssim$ relation from §4.4, several other WP combinators are derived to be used internally by the F^* type-checker; again previously these had to be written by hand.

5. EMF^* with primitive state

As we have seen in §3, EMF^* encodes all its effects using pure functions. However, one would like to be able to run F^* programs efficiently using primitive effects. In this section, we show how EMF^* ’s pure monads apply to F^* ’s existing compilation strategy, which provides primitive support for state via compilation to OCaml, which, of course, has state natively.⁷ The main theorem of §5.1 states that well-typed EMF^* programs using the state monad abstractly (i.e., not breaking the abstraction of the state monad with arbitrary uses of reify and reflect) are related by a simulation to EMF_{ST}^* programs that execute with a primitive notion of state. This result exposes a basic tension: although very useful for proofs, reify and reflect can break the abstractions needed for efficient compilation. In §5.2, we show how to get the best of both worlds by relying on F^* ’s Ghost effect—by restricting reify and reflect to computationally irrelevant code (e.g., proofs), we can use them freely for proving and then erase them before compilation for efficient execution.

5.1 EMF_{ST}^* : A sub-language of EMF^* with primitive state

The syntax of EMF_{ST}^* corresponds to EMF^* , except, we configure it to just use the ST monad. Other effects that may be added to EMF^* can already be expanded into their encodings in its primitive Pure monad—as such, we think of EMF_{ST}^* as modeling a compiler target for EMF^* programs extended with ST implemented primitively, and other arbitrary effects implemented purely. We thus exclude from EMF_{ST}^* the reify and reflect operators and drop type and WP arguments of return, bind and lift operators, since these are no longer relevant here.

The operational semantics of EMF_{ST}^* is a small-step, call-by-value reduction relation between pairs (s, e) of a state s and a term e . The relation includes the pure reduction steps of EMF^* simply carrying the state along (we only show ST-beta), and three primitive reduction rules for ST, shown below. The only irreducible ST computation is $\text{ST.return } v$. The term $\text{ST.bind } e \text{ x.e'}$ reduces even without an enclosing reify, since the state is primitive.

$$\begin{aligned} (s, (\lambda x.t.e)v) &\rightsquigarrow (s, e[v/x]) & \text{ST-beta} \\ (s, \text{ST.bind } (\text{ST.return } v) \text{ x.e}) &\rightsquigarrow (s, e[v/x]) & \text{ST-bind} \\ (s, \text{ST.get } ()) &\rightsquigarrow (s, \text{ST.return } s) & \text{ST-get} \\ (s, \text{ST.put } s') &\rightsquigarrow (s', \text{ST.return } ()) & \text{ST-put} \end{aligned}$$

Relating EMF^* to EMF_{ST}^* We define a (partial) translation from EMF^* to EMF_{ST}^* , and show that one or more steps of reduction in EMF_{ST}^* are matched by one or more steps in EMF^* . This result guarantees that it is sound to verify a program in EMF^* and execute it in EMF_{ST}^* , since the verification holds for all EMF^* reduction sequences, and EMF_{ST}^* evaluation corresponds to one such reduction.

The main intuition behind our proof is that the reduction of reflect-free EMF^* programs maintains terms in a very specific structure—a stateful redex reduces in a context structured like

⁷ F^* also compiles exceptions natively to OCaml, however we focus only on state here, leaving a formalization of primitive exceptions to the future—we expect it to be similar to the development here.

a telescope of binds, with the state threaded sequentially as the telescope evolves. We describe this invariant structure as an EMF^* context, K , parameterized by a state s , where $\hat{E}$ is a single-hole, reify-and-reflect-free EMF^* context, a refinement of the evaluation contexts of §3, to be filled by a reify-and-reflect free EMF^* term, f . Additionally, we separate the $\hat{E}$ contexts by their effect into several sorts; $\hat{E} : \text{Tot}$ and $\hat{E} : \text{Pure}$ are contexts which when filled by a suitably typed term produce in EMF^* a Tot or Pure term, respectively; the case $\hat{E} : \text{Inert}$ is for an un-reified stateful EMF^* term. The last two cases are the most interesting, representing the base and inductive case of the telescope of a stateful term “caught in the act” of reducing—we refer to them as the Active contexts. We omit the sort of a context when it is irrelevant.

$K s ::= \hat{E} : \text{Tot} \mid \hat{E} : \text{Pure} \mid \hat{E} : \text{Inert} \mid \text{reify } \hat{E} s : \text{Active}$
 $\mid \text{Pure.bind } (K s) p.((\lambda x.\text{reify } f) (\text{fst } p) (\text{snd } p)) : \text{Active}$

Next, we define a simple translation $\llbracket \cdot \rrbracket$ from contexts Ks to EMF_{ST}^* .

$\llbracket \hat{E} \rrbracket = \hat{E}$
 $\llbracket \text{reify } \hat{E} s \rrbracket = \hat{E}$
 $\llbracket \text{Pure.bind } (K s) p.((\lambda x.\text{reify } f) (\text{fst } p) (\text{snd } p)) \rrbracket$
 $= \text{ST.bind } \llbracket K s \rrbracket x.f$

Theorem 10 (Simulation). *For all well-typed closed, filled contexts $K s f$, either $K s$ is Inert, or one of the following is true:*

- (1) $\exists K' s' f'. (s, \llbracket K s \rrbracket f) \rightsquigarrow^+ (s', \llbracket K' s' \rrbracket f')$
and $K s f \longrightarrow^+ K' s' f'$ and $\text{sort } (K s) = \text{sort } (K' s')$
and if $K' s'$ is not Active then $s = s'$.
- (2) $K s$ is Active and $\exists v s'. (s, \llbracket K s \rrbracket f) \rightsquigarrow^* (s', \text{ST.return } v)$
and $K s f \longrightarrow^+ \text{Pure.return } (v, s')$.
- (3) $K s$ is Pure and $\exists v. \llbracket K s \rrbracket f = K s f = \text{Pure.return } v$.
- (4) $K s$ is Tot and $\exists v. \llbracket K s \rrbracket f = K s f = v$.

5.2 Restoring reify and reflect for extrinsic proofs

Theorem 10 applies only to programs that lack reflect and use reify only in a very specific manner, as described by the telescoping invariant. To apply the theorem to F^* , we make use of its facilities for erasing computationally irrelevant code and its module system. To justify F^* ’s primitive effect implementation, we argue that F^* programs using the ST monad are free of reify and reflect, *after erasure* of computationally irrelevant code, and hence the post-erasure programs are candidates for Theorem 10. We briefly describe how this works, next.

First, as described in §2.5, we use the module system to hide both the reify and reflect operators from clients of a module $F\text{Star.State}$ defining the ST effect. We expose to clients only `ghost_reify`, a function equivalent to reify, but exposed to clients at the signature shown below. Notice that the function’s co-domain is marked with the Ghost effect, meaning that it can only be used within specifications (e.g., WPs and assertions)—any other use will be flagged as a typing error by F^* .

`ghost_reify: ST a wp  $\rightarrow$  Ghost (s0:s  $\rightarrow$  Pure (a * s) (wp s0))`

Next, within $F\text{Star.State}$, we implement the following combinator, `refine_ST`, a total function that allows a client to augment the specification of an effectful function f from some wp to a weaker one that additionally records that the value returned by f on any argument and input state $s0$ corresponds to the computational behavior of the (ghostly) reification of f .

`refine_ST: f:(x:a  $\rightarrow$  ST b (wp x))`
 `$\rightarrow$  Tot (x:a  $\rightarrow$  ST b ( $\lambda$  s0 post  $\rightarrow$`
`wp x s0 ( $\lambda$  r  $\rightarrow$  r = ghost_reify (f x) s0  $\implies$  post r)))`

6. Related work

We have already discussed many elements of related work throughout the paper. Here we focus on a few themes not covered fully elsewhere.

Representing monads Our work draws a lot from Filinski’s (1994) monadic reflection methodology, for representing and controlling the abstraction of monads. In particular, our DM monad definition language is essentially the language of (Filinski 1994) with some restrictions on the allowed types. Beyond controlling abstraction, Filinski shows how monadic reflection enables a universal implementation of monads using composable continuations and a single mutable cell. We do not (yet) make use of that aspect of his work, partly because deploying this technique in practice is challenging, since it requires compiling programs to a runtime system that provides composable continuations. Filinski’s (1999) work on representing layered monads generalizes his technique to the setting of multiple monads. We also support multiple monads, but instead of layering monads, we define each monad purely, and relate them via morphisms. This style is better suited to our purpose, since one of our primary uses of reification is purification, i.e., revealing the pure representation of an effectful term for reasoning purposes. With layering, multiple steps of reification may be necessary, which may be inconvenient for purification. Finally, Filinski (2010) gives an operational semantics that is extensible with monadic actions, taking the view of effects as being primitive, rather than encoded purely. We take a related, but slightly different view: although effects are encoded purely in EMF^* , we see it as language in which to analyze and describe the semantics of a primitively effectful object language, EMF_{ST}^* , relating the two via a simulation.

Dependent types and effects Nanevski et al. developed Hoare type theory (HTT) (Nanevski et al. 2008) and Ynot (Chlipala et al. 2009) as a way of extending Coq with effects. The strategy there is to provide an axiomatic extension of Coq with a single catch-all monad in which to encapsulate imperative code. Being axiomatic, their approach lacks the ability to reason extrinsically about effectful terms by computation. However, their approach accommodates effects like non-termination, which EMF^* currently lacks. Interestingly, the internal semantics of HTT is given using predicate transformers, similar in spirit to EMF^* ’s WP semantics. It would be interesting to explore whether or not our free proofs of monotonicity and conjunctivity simplify the proof burden on HTT’s semantics.

Zombie (Casinghino et al. 2014) is a dependently typed language with general recursion, which supports reasoning extrinsically about potentially divergent code—this approach may be fruitful to apply to EMF^* to extend its extrinsic reasoning to divergent code.

Another point in the spectrum between extrinsic and intrinsic reasoning is Charguéraud’s (2011) characteristic formulae, which provide a precise formula in higher-order logic capturing the semantics of a term, similar in spirit to our WPs. However, as opposed to WPs, characteristic formulae are used interactively to prove program properties after definition, although not via computation, but via logical reasoning. Interesting enough, characteristic formulae are structured in a way that almost gives the illusion that they are the terms themselves. CFML is tool in Coq based on these ideas, providing special tactics to manipulate formulas structured this way.

Brady (2013) encodes algebraic effects with pre- and postconditions in Idris in the style of Atkey’s (2009) parameterized monads. Rather than speaking about the computations themselves, the pre- and postconditions refer to some implicit state of the world, e.g., whether or not a file is closed. In contrast, F^* ’s WPs give a full logical characterization of a computation. Additionally, the WP style is better suited to computing verification conditions, instead of explicitly chaining indices in the parameterized monad.

It would be interesting, and possibly clarifying, to link up with recent work on the denotational semantics of effectful languages with dependent types (Ahman et al. 2016); in our case one would investigate the semantics of EMF^* and EMF_{ST}^* , which has state, but extended with recursion (and so with nontermination).

Continuations and predicate transformers We are not the first to study the connection between continuations and predicate transformers. For example, Jensen (1978) and Audebaud and Zucca (1999) both derive WPs from a continuation semantics of first-order imperative programs. While they only consider several primitive effects, we allow arbitrary monadic definitions of effects. Also while their work is limited to the first-order case, we formalize the connection between WPs and CPS also for higher-order. The connection between WPs and the continuation monad also appears in Keimel (2015); Keimel and Plotkin (2016).

7. Looking back, looking ahead

While our work has yielded the pleasant combination of both a significant simplification and boost in expressiveness for F^* , we believe it can also provide a useful foundation on which to add user-defined effects to other dependently typed languages. All that is required is the Pure monad upon which everything else can be built, mostly for free.

On the practical side, going forward, we hope to make use of the new extrinsic proving capabilities in F^* to simplify specifications and proofs in several ongoing program verification efforts that use F^* . We are particularly interested in furthering the relational verification style, sketched in §2.6. We also hope to scale EMF^* to be a definitive semantics of all of F^* —the main missing ingredients are recursion and its semantic termination check, inductive types, universe polymorphism, and the extensional treatment of equality.

Along another axis, we have already mentioned our plans to add non-termination (§2.4) to DM and EMF^* , and to investigate translations of effect handlers (§2.5). We also hope to enhance DM in other ways, e.g., relaxing the stratification of types and adding inductive types which will allow us to define monads for some forms of nondeterminism and many forms of I/O. Enriching DM further, one could also add dependent types, reducing the gap between it and F^* , and bringing within reach examples like Ahman and Uustalu’s (2013) dependently typed update monads.

Acknowledgments We are grateful to Pierre-Evariste Dagand, Michael Hicks, and Kenji Maillard for interesting discussions. We also thank the anonymous reviewers for their helpful feedback.

A. Appendix

In this appendix we provide proofs and auxiliary results for the theorems that appear in the body of the paper. We also show the full type system for the source language.

A.1 The definitional language DM

In the typing judgment, the metavariable Δ represents a set of type variables that remains fixed throughout typing. It is used to introduce top-level let-polymorphism on all CPS'd/elaborated terms. A type is well-formed in the context Δ if all of its variables are in Δ . In rigor, all judgments from here onwards are subject to that constraint, which we do not write down. A context Γ is well-formed if both (1) all of its types are well-formed according to Δ (2) no variable names are repeated. This last condition simplifies reasoning about substitution and does not limit the language in any way.

We assume that every base type in DM is also a base type in EMF* (or that there exists a mapping from them, formally), and that source constants are also present and with the same type (formally, also a mapping for constants that respects the previous one).

The typing judgment for DM is given in Figure 6. We assume that the types appearing in the rules are well-formed. For example, in the (ST-PAIR) rule, either both H and H' are in A or both are in C etc.

A.2 CPS translation (WP generation)

The full $\star$ -translation for DM expressions is given in Figure 7. The one for types was previously defined. We define translation on environments in the following way:

$$\frac{\Delta = X_1, \dots, X_m}{\Delta^* = X_1 : \text{Type}_0, \dots, X_m : \text{Type}_0} \quad \frac{\Gamma = x_1 : H_1, \dots, x_n : H_n}{\Gamma^* = x_1 : H_1^*, \dots, x_n : H_n^*}$$

One can then prove the following:

Lemma 11 (Well-typing of $\star$ -translation). *For any Γ, e, A and H :*

$$\begin{aligned} \Delta \mid \Gamma \vdash e : H ! n &\implies \Delta^*, \Gamma^* \vdash e^* : H^* \\ \Delta \mid \Gamma \vdash e : A ! \tau &\implies \Delta^*, \Gamma^* \vdash e^* : (A^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0 \end{aligned}$$

Proof. By induction on the typing derivation. $\square$

In this lemma statement, and in those that follow, when writing e^* we refer to the translation of e using the typing derivation from the premise.

A.3 Elaboration

The definitions of $\underline{A}$ and the F relation were previously given. For elaboration, we also translate environments, in the following manner:

$$\frac{\Delta = X_1, \dots, X_m}{\underline{\Delta} = X_1 : \text{Type}_0, \dots, X_m : \text{Type}_0}$$

$$\frac{\Gamma = x_1 : H_1, \dots, x_n : H_n}{\underline{\Gamma} = x_1 : H_1, \dots, x_n : H_n}$$

Note that for any computational variable in the context, we introduce two variables: one for its WP and one for its actual expression. The x^w variable, which is assumed to be fresh, is used only at the WP level. Also note that $\underline{\Delta} = \Delta^*$.

For any Γ , we define the substitution s_Γ as $[x_{i_1}^w/x_{i_1}, \dots, x_{i_k}^w/x_{i_k}]$, for the computational variables $x_{i_1}, \dots, x_{i_k} \in \Gamma$.

Similarly to Lemma 11 we show that:

Lemma 12 (Well-typing of $\star$ -translation — elaboration contexts). *For any Γ, e, A and C we have:*

$$\begin{aligned} \Delta \mid \Gamma \vdash e : H ! n &\implies \underline{\Delta}, \underline{\Gamma} \vdash e^* s_\Gamma : H^* \\ \Delta \mid \Gamma \vdash e : A ! \tau &\implies \underline{\Delta}, \underline{\Gamma} \vdash e^* s_\Gamma : (A^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0 \end{aligned}$$

For expression elaboration we aim to show that:

$$\frac{\Delta \mid \Gamma \vdash e : A ! n}{\underline{\Delta}, \underline{\Gamma} \vdash e : \underline{A}} \quad \frac{\Delta \mid \Gamma \vdash e : C ! n}{\underline{\Delta}, \underline{\Gamma} \vdash e : F_C e^* s_\Gamma} \quad \frac{\Delta \mid \Gamma \vdash e : A ! \tau}{\underline{\Delta}, \underline{\Gamma} \vdash e : \text{Pure } \underline{A} (e^* s_\Gamma)}$$

A.4 Statement of the logical relations lemma

Our main theorem is the fact the $\star$ -translation of a term is properly related, via the F relation, to its elaboration. This means it provides an adequate logical representation of the term.

Theorem 13 (Logical relations lemma). *For any Δ, Γ, e, C, A :*

$$\begin{aligned} \Delta \mid \Gamma \vdash e : C ! n &\implies \underline{\Delta}, \underline{\Gamma} \vdash e : F_C e^* s_\Gamma \\ \Delta \mid \Gamma \vdash e : A ! \tau &\implies \underline{\Delta}, \underline{\Gamma} \vdash e : \text{Pure } \underline{A} (e^* s_\Gamma) \end{aligned}$$

Note that in both cases, $e : H ! \varepsilon$ implies $e : G_H^\varepsilon(e^ s_\Gamma)$, as per the definition of G given before.*

As a trivial corollary of this theorem, we get that if a computational term is typable in the empty context ($\Delta \mid \cdot \vdash e : C ! n$) then $\underline{\Delta} \vdash e : F_C e^*$, as stated in the paper.

A.5 Proof of the logical relation lemma

Theorem 14. *For any $A, \Delta \mid \Gamma \vdash e : A ! n \implies \underline{e} = e^* s_\Gamma$. (That is, syntactic equality).*

Proof. By induction on the typing derivation. The cases for (ST-RET) and (ST-BIND) do not apply.

(1) (ST-VAR)

Our goal is to show $x = x s_\Gamma$. Since the type of x is A the substitution does not affect x , thus they're trivially both x .

(2) (ST-CONST)

Say $\Delta \mid \Gamma \vdash \kappa(b_1, \dots, b_n) : b ! n$. By the induction hypothesis we know that $\underline{b_i} = b_i^* s_\Gamma$ for each i . We thus trivially get our goal by substitution of the arguments.

(3) (ST-ABS)

Say we concluded $\Delta \mid \Gamma \vdash \lambda x : A. e : A \xrightarrow{n} A' ! n$. Our premise is (note the substitution from the IH does not affect x , as it has an A -type) the fact that $\underline{e} = e^* s_\Gamma$. We need to show that:

$$\underline{\lambda x : A. e} = (\lambda x : A. e)^*$$

which is just

$$\lambda x : \underline{A}. \underline{e} = \lambda x : A^*. e^*$$

which is trivial from our hypothesis and since $\underline{A} =_{\text{def}} A^*$.

(4) (ST-APP)

Say we concluded $\Delta \mid \Gamma \vdash f e : A' ! n$ by the premises

$$\Delta \mid \Gamma \vdash f : A \xrightarrow{n} A' ! n \quad \Delta \mid \Gamma \vdash e : A ! n$$

(it cannot be the case that e has some C type, because of the type restrictions). From the inductive hypotheses we have:

$$\underline{f} = f^* s_\Gamma \implies \underline{f e} = (f^* s_\Gamma) \underline{e} \implies$$

$$\underline{f e} = (f^* s_\Gamma) (e^* s_\Gamma) \implies \underline{f e} = (f e)^* s_\Gamma$$

As required.

(5) (ST-FST), (ST-SND), (ST-PAIR), (ST-INL), (ST-INR)

All of these are trivial by applying the IH. For (ST-PAIR) one needs to note that the restrictions will ensure that the type of the pair will be an A -type.

$\frac{\text{ST-VAR} \quad x : H \in \Gamma}{\Delta \mid \Gamma \vdash x : H!n}$	$\frac{\text{ST-CONST} \quad \Delta \mid \Gamma \vdash e_i : b_i!n \quad \kappa : b_1, \dots, b_n \rightarrow b}{\Delta \mid \Gamma \vdash \kappa(e_1, \dots, e_n) : b!n}$	$\frac{\text{ST-ABS} \quad \Delta \mid \Gamma, x : H \vdash e : H'! \varepsilon}{\Delta \mid \Gamma \vdash \lambda x : H. e : H \xrightarrow{\varepsilon} H'!n}$
$\frac{\text{ST-APP} \quad \Delta \mid \Gamma \vdash e : H \xrightarrow{\varepsilon} H'!n \quad \Delta \mid \Gamma \vdash e' : H!n}{\Delta \mid \Gamma \vdash ee' : H'! \varepsilon}$	$\frac{\text{ST-PAIR} \quad \Delta \mid \Gamma \vdash e : H!n \quad \Delta \mid \Gamma \vdash e' : H'!n}{\Delta \mid \Gamma \vdash (e, e') : H \times H'!n}$	$\frac{\text{ST-FST} \quad \Delta \mid \Gamma \vdash e : H \times H'!n}{\Delta \mid \Gamma \vdash \text{fst}(e) : H!n}$
$\frac{\text{ST-INL} \quad \Delta \mid \Gamma \vdash e : A!n}{\Delta \mid \Gamma \vdash \text{inl}(e) : A + A'!n}$	$\frac{\text{ST-CASE} \quad \Delta \mid \Gamma \vdash e : A + A'!n \quad \Delta \mid \Gamma, x : A \vdash e_1 : H! \varepsilon \quad \Delta \mid \Gamma, x : A' \vdash e_2 : H! \varepsilon}{\Delta \mid \Gamma \vdash \text{case } e \text{ inl } x : A. e_1; \text{inr } y : A'. e_2 : H! \varepsilon}$	
$\frac{\text{ST-RET} \quad \Delta \mid \Gamma \vdash e : A!n}{\Delta \mid \Gamma \vdash \text{return}_\tau e : A! \tau}$	$\frac{\text{ST-BIND} \quad \Delta \mid \Gamma \vdash e : A! \tau \quad \Delta \mid \Gamma, x : A \vdash e' : A'! \tau}{\Delta \mid \Gamma \vdash \text{bind}_\tau e \text{ to } x : A \text{ in } e' : A'! \tau}$	

Figure 6. Typing rules of DM

x^*	$= x$	$K(e_1, \dots, e_n)^*$	$= K e_1^* \dots e_n^*$
$(f e)^*$	$= f^* e^*$	$(\lambda x : H. e)^*$	$= \lambda x : H^*. e^*$
$\text{fst}(e)^*$	$= \text{fst } e^*$	$\text{snd}(e)^*$	$= \text{snd } e^*$
$\text{inl}(e)^*$	$= \text{inl } e^*$	$\text{inr}(e)^*$	$= \text{inr } e^*$
$(e_1, e_2)^*$	$= (e_1^*, e_2^*)$	$(\text{case } e_0 \text{ inl } x : A. e_1; \text{inr } y : A'. e_2)^*$	$= \text{case}(e_0^*) x. e_1^* y. e_2^*$
$(\text{return}_\tau e)^*$	$= \lambda p : A^* \rightarrow \text{Type}_0. p e^*$	$(\text{when } \Delta \mid \Gamma \vdash e : A!n)$	
$(\text{bind}_\tau e_1 \text{ to } x \text{ in } e_2)^*$	$= \lambda p : A'^* \rightarrow \text{Type}_0. e_1^* (\lambda x : A. e_2^* p)$	$(\text{when } \Delta \mid \Gamma, x : A \vdash e_2 : A'! \tau)$	

Figure 7. Definition of the $\star$ -translation for DM terms

(6) (ST-CASE)

Say we concluded $\Delta \mid \Gamma \vdash \text{case } e \text{ inl } x : A_0. e_1; \text{inr } y : A_1. e_2 : A_2!n$. As inductive hypothesis we have:

$$\underline{e} = e^* \text{ s}_\Gamma \quad \underline{e}_1 = e_1^* \text{ s}_\Gamma \quad \underline{e}_2 = e_2^* \text{ s}_\Gamma$$

(e_1 and e_2 are typed in the context extended with x and y respectively, however since they are A -types the substitution is the same)

The goal is:

$$\begin{aligned} & (\text{case}(\underline{e}) x. \underline{e}_1 y. \underline{e}_2) \\ &= (\text{case}(e^* \text{ s}_\Gamma) x. e_1^* \text{ s}_\Gamma y. e_2^* \text{ s}_\Gamma) \end{aligned}$$

From the IHs, and since $\underline{A} =_{\text{def}} A^*$ we get our goal.

□

(3) (ST-FST), (ST-SND), (ST-PAIR), (ST-INL), (ST-INR)

Trivial by using IH.

(4) (ST-CASE)

Say we concluded $\Delta \mid \Gamma \vdash \text{case } e \text{ inl } x : A_0. e_1; \text{inr } y : A_1. e_2 : A_2!n$ by (ST-CASE). Our IHs give us

$$\begin{aligned} \underline{\Delta}, \underline{\Gamma} &\vdash \underline{e} : \underline{A_0} + \underline{A_1} \\ \underline{\Delta}, \underline{\Gamma}, x : \underline{A_0} &\vdash \underline{e_1} : \underline{A_2} \\ \underline{\Delta}, \underline{\Gamma}, y : \underline{A_1} &\vdash \underline{e_2} : \underline{A_2} \end{aligned}$$

By a non-dependent application of T-CaseTot we get

$$\underline{\Delta}, \underline{\Gamma} \vdash \text{case}(\underline{e}) x. \underline{e_1} y. \underline{e_2} : \underline{A_2}$$

Which is our goal.

(5) (ST-ABS), (ST-APP)

Both trivial from IHs.

□

Theorem 15. *If $\Delta \mid \Gamma \vdash e : A!n$, then $\underline{\Delta}, \underline{\Gamma} \vdash \underline{e} : \underline{A}$.*

Proof. By induction on the typing derivation.

(1) (ST-VAR)

We have $\Delta \mid \Gamma \vdash x : A!n$, with $x \in \Gamma$. By the translation for environments, we have $x : \underline{A}$ in $\underline{\Gamma}$, so this is trivial.

(2) (ST-CONST)

For any constant $\kappa : (b_1, \dots, b_n) \rightarrow b$ say we have $\Delta \mid \Gamma \vdash \kappa(e_1, \dots, e_n) : b!n$ by (ST-CONST) (note that b and all the b_i are in A). This means that for every i we have as inductive hypothesis:

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{e_i} : \underline{b_i}$$

Since κ is also a target constant of the same type, we thus have:

$$\underline{\Delta}, \underline{\Gamma} \vdash \kappa \underline{e_1} \dots \underline{e_n} : \underline{b}$$

Which is exactly our goal as $\underline{b} = b$.

Before jumping into the logical relation lemma, we will require the following auxiliary lemma, of which we make heavy use.

Lemma 16 (Invariancy of $F_C w$). *If $\Gamma \models w_1 = w_2$, then $\Gamma \vdash F_C w_1 < F_C w_2$.*

Proof. By induction on C .

(1) $C \xrightarrow{\tau} A$

We need to show that

$$\Gamma \vdash F_{C \xrightarrow{\tau} A} w_1 < F_{C \xrightarrow{\tau} A} w_2$$

Which is

$$\Gamma \vdash x^w : C^* \rightarrow F_C x^w \rightarrow \text{Pure } \underline{A} (w_1 x^w) < x^w : C^* \rightarrow F_C x^w \rightarrow \text{Pure } \underline{A} (w_2 x^w)$$

After two applications of (ST-PROD) (and some trivial reflexivity discharges), the required premise to show is:

$$\Gamma, x^w : C^*; F_C x^w \vdash \text{Pure } \underline{A} (w_1 x^w) < \text{Pure } \underline{A} (w_2 x^w)$$

By (S-PURE) we're required to show that $\underline{A}$ is a subtype of itself (which is trivial by reflexivity of subtyping (S-CONV)) and that w_2 is stronger than w_1 , which can be easily proven as they are equal.

(2) $A \xrightarrow{\tau} A$

Very similar to the previous case, but simpler.

(3) $C \xrightarrow{n} C'$

We need to show that

$$\Gamma \vdash F_{C \xrightarrow{n} C'} w_1 <: F_{C \xrightarrow{n} C'} w_2$$

Which is

$$\Gamma \vdash x^w : C^* \rightarrow F_C x^w \rightarrow F_{C'} w_1 x^w <: x^w : C^* \rightarrow F_C x^w \rightarrow F_{C'} w_1 x^w$$

After two applications of (ST-PROD) (and some trivial reflexivity discharges), the required premise to show is:

$$\Gamma, x^w : C^*, F_C x^w \vdash F_{C'} w_1 x^w <: F_{C'} w_2 x^w$$

As in this context we can show $w_1 x^w = w_2 x^w$ we apply our IH to the type C' and are done.

(4) $A \xrightarrow{n} C'$

Also very similar to the previous case, but simpler.

(5) $C \times C'$

Trivial by IHs and concluding that $\text{fst } w_1 = \text{fst } w_2$, and similarly for snd .

□

Proof of Theorem 13 (Logical relations lemma)

Proof. The two parts are proved by a joint structural induction.

(1) (ST-VAR)

We have $\Delta \mid \Gamma \vdash x : C!n$, with $x \in \Gamma$. By the translation for environments, we have $x^w : C^*$ and $x : F_C x^w$ in $\underline{\Gamma}$. Since x is covered by the substitution s_Γ , what we need to prove is $\underline{\Delta}, \underline{\Gamma} \vdash x : F_C x^w$, which is exactly what we have in the environment.

(2) (ST-PAIR)

Suppose we proved $(e_1, e_2) : C_1 \times C_2!n$ by (ST-PAIR). We want to show: $\underline{\Delta}, \underline{\Gamma} \vdash (\underline{e}_1, \underline{e}_2) : F_{C_1 \times C_2} (e_1^*, e_2^*) s_\Gamma$, i.e., that (after reduction inside F):

$$\underline{\Delta}, \underline{\Gamma} \vdash (\underline{e}_1, \underline{e}_2) : F_{C_1} e_1^* s_\Gamma \times F_{C_2} e_2^* s_\Gamma$$

This is trivial by applying both IHs.

(3) (ST-FST), (ST-SND)

Suppose we proved $\Delta \mid \Gamma \vdash \text{fst}(e) : C_1!n$ by (ST-FST). We need to then show

$$\underline{\Delta}, \underline{\Gamma} \vdash \text{fst } \underline{e} : F_{C_1} \text{fst } e^* s_\Gamma$$

By our induction hypothesis we have $\underline{\Delta}, \underline{\Gamma} \vdash \underline{e} : F_{C_1 \times C_2} e^* s_\Gamma$, which is

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{e} : F_{C_1} \text{fst } e^* s_\Gamma \times F_{C_2} \text{snd } e^* s_\Gamma$$

It is therefore easy to see that we have our goal.

(4) (ST-ABS)

There are two cases:

• $A \xrightarrow{\varepsilon} H$

Suppose we concluded $\Delta \mid \Gamma \vdash \lambda x : A. e : A \xrightarrow{\varepsilon} H!n$. Then we have $\Delta \mid \Gamma, x : A \vdash e : H! \varepsilon$ and so, by the induction hypothesis, in both cases for ε we have

$$\underline{\Delta}, \underline{\Gamma}, x : \underline{A} \vdash \underline{e} : G_H^\varepsilon(e^* s_\Gamma)$$

And we have to show:

$$\underline{\Delta}, \underline{\Gamma} \vdash \lambda x : \underline{A}. \underline{e} : F_{A \xrightarrow{n} H} (\lambda x : A^*. e^*) s_\Gamma$$

Which is

$$\underline{\Delta}, \underline{\Gamma} \vdash \lambda x : \underline{A}. \underline{e} : x : \underline{A} \rightarrow G_H^\varepsilon((\lambda x : A^*. e^*) s_\Gamma x)$$

Since the substitution does not cover x , the argument to G is just $e^* s_\Gamma$, thus we use our IH to conclude this easily.

• $C \xrightarrow{\varepsilon} H$

Suppose we concluded $\Delta \mid \Gamma \vdash \lambda x : C. e : C \xrightarrow{\varepsilon} H!n$. Then we have $\Delta \mid \Gamma, x : C \vdash e : H! \varepsilon$ and so, by the induction hypothesis we have, in either case for ε :

$$\underline{\Delta}, \underline{\Gamma}, x^w : C^*, x : F_C x^w \vdash \underline{e} : G_H^\varepsilon(e^* s_\Gamma [x^w/x])$$

And we have to show:

$$\underline{\Delta}, \underline{\Gamma} \vdash \lambda x^w : C^*. \lambda x : F_C x^w. \underline{e} : F_{C \xrightarrow{\varepsilon} H} (\lambda x : C^*. e^*) s_\Gamma$$

Which is

$$\begin{aligned} \underline{\Delta}, \underline{\Gamma} \vdash \lambda x^w : C^*. \lambda x : F_C x^w. \underline{e} \\ : x^w : C^* \rightarrow F_C x^w \rightarrow G_H^\varepsilon((\lambda x : C^*. e^*) s_\Gamma x^w) \end{aligned}$$

Using T-Abs twice we can conclude this via

$$\underline{\Delta}, \underline{\Gamma}, x^w : C^*, x : F_C x^w \vdash \underline{e} : G_H^\varepsilon((\lambda x : C^*. e^*) s_\Gamma x^w)$$

Since the substitution does not cover x , the argument to F reduces to $e^* s_\Gamma [x^w/x]$, thus we use our IH to conclude this easily.

(5) (ST-APP)

Again, There are two possible cases:

• $A \xrightarrow{\varepsilon} H$

We concluded $\Delta \mid \Gamma \vdash fe : G! \varepsilon$. Our premises are $\Delta \mid \Gamma \vdash f : A \xrightarrow{\varepsilon} C!n$ and $\Delta \mid \Gamma \vdash e : A!n$. The IH for f is, expanding F:

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{f} : x : \underline{A} \rightarrow G_H^\varepsilon((f^* s_\Gamma) x)$$

By T-App, and since $\underline{e} : \underline{A}$, this is just:

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{f} \underline{e} : G_H^\varepsilon((f^* s_\Gamma) \underline{e})$$

Since from a previous theorem we know we have $\underline{e} = e^* s_\Gamma$ (syntactically), we can conclude:

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{f} \underline{e} : G_H^\varepsilon((f^* s_\Gamma) (e^* s_\Gamma))$$

This is exactly:

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{f} \underline{e} : G_H^\varepsilon((f e)^* s_\Gamma)$$

which is our goal, in either the $C!n$ or the $A! \tau$ case.

• $C \xrightarrow{\varepsilon} H$

We concluded $\Delta \mid \Gamma \vdash fe : H! \varepsilon$. Our premises are $\Delta \mid \Gamma \vdash f : C \xrightarrow{\varepsilon} H!n$ and $\Delta \mid \Gamma \vdash e : C!n$. The IHs are, expanding F:

$$\begin{aligned} \underline{\Delta}, \underline{\Gamma} \vdash \underline{f} : x^w : C^* \rightarrow F_C x^w \rightarrow G_H^\varepsilon((f^* s_\Gamma) x^w) \\ \underline{\Delta}, \underline{\Gamma} \vdash \underline{e} : F_C e^* s_\Gamma \end{aligned}$$

Thus by two uses of T-App (noting that it's well typed by our IH for e), we can conclude:

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{f} (e^* s_\Gamma) \underline{e} : G_{C'}^\varepsilon((f^* s_\Gamma) (e^* s_\Gamma))$$

This is just, syntactically:

$$\underline{\Delta}, \underline{\Gamma} \vdash \underline{f} \underline{e} : G_{C'}^\varepsilon((f^* s_\Gamma) (e^* s_\Gamma))$$

which is our goal, in either the $C!n$ or the $A! \tau$ case.

(6) (ST-CASE)

There are two cases depending on whether we eliminate into $C!n$ or $A! \tau$. Both of these cases are quite dull, and deal mostly with the typing judgment on the target. This may be skipped without hindering any of the main ideas.

• $C!n$

Suppose that $\Delta \mid \Gamma \vdash e : A + A'!n$, $\Delta \mid \Gamma, x : A \vdash e_1 : C!n$, and $\Delta \mid \Gamma, y : A' \vdash e_2 : C!n$, so that $\Delta \mid \Gamma \vdash \text{case } e \text{ in } x : A. e_1; \text{ in } y : A'. e_2 : C!n$. We will go into detail only for e_1

as the typing and reasoning for e_2 is exactly analogous. As inductive hypothesis for e_1 we have $\square$

$$\Delta, \Gamma, x : \underline{A} \vdash \underline{e}_1 : F_C e_1^* s_\Gamma$$

We wish to show:

$$\begin{aligned} \Delta, \Gamma \vdash \text{case}_{F_C} \text{case}(z) e_1^* s_\Gamma e_2^* s_\Gamma (\underline{e} \text{ as } z) x : \underline{A} \cdot \underline{e}_1 y : \underline{A}' \cdot \underline{e}_2 \\ : F_C (\text{case}(e^*) x \cdot e_1^* y \cdot e_2^*) s_\Gamma \end{aligned}$$

Which is

$$\begin{aligned} \Delta, \Gamma \vdash \text{case}_{F_C} \text{case}(z) x \cdot e_1^* s_\Gamma y \cdot e_2^* s_\Gamma (\underline{e} \text{ as } z) x : \underline{A} \cdot \underline{e}_1 y : \underline{A}' \cdot \underline{e}_2 \\ : F_C \text{case}(e^* s_\Gamma) x \cdot e_1^* s_\Gamma y \cdot e_2^* s_\Gamma \end{aligned}$$

Since $\underline{e} = e^* s_\Gamma$ we will prove this has type $F_C \text{case}(\underline{e}) x \cdot e_1^* s_\Gamma y \cdot e_2^* s_\Gamma$.

By T-CaseTot, we should show:

$$\begin{aligned} \Delta, \Gamma \vdash \underline{e} : \underline{A} + \underline{A}' \\ \Delta, \Gamma, x : \underline{A} \vdash \underline{e}_1 : F_C \text{case}(\text{inl } x) x \cdot e_1^* s_\Gamma y \cdot e_2^* s_\Gamma \end{aligned}$$

(And the one for e_2). We get the first one trivially by theorem 15. The second is by reduction equivalent to:

$$\Delta, \Gamma, x : \underline{A} \vdash \underline{e}_1 : F_C e_1^* s_\Gamma$$

Which is exactly our IH for e_1 , so we're done.

- $A! \tau$

Our hypotheses are:

$$\begin{aligned} \Delta, \Gamma \vdash \underline{e} : A_0 + A_1 \\ \Delta, \Gamma, x : \underline{A}_0 \vdash \underline{e}_1 : \text{Pure } \underline{A}_2 (e_1^* s_\Gamma) \\ \Delta, \Gamma, y : \underline{A}_1 \vdash \underline{e}_2 : \text{Pure } \underline{A}_2 (e_2^* s_\Gamma) \end{aligned}$$

Applying T-Case (non-dependently) we get.

$$\begin{aligned} \Delta, \Gamma \vdash \text{case}(\underline{e}) x \cdot \underline{e}_1 y \cdot \underline{e}_2 : \\ \text{Pure } \underline{A}_2 (\text{case}(\underline{e}) x \cdot e_1^* s_\Gamma y \cdot e_2^* s_\Gamma) \end{aligned}$$

Since we know $\underline{e} = e^* s_\Gamma$ and since $\underline{A} =_{\text{def}} A^*$ this is exactly:

$$\begin{aligned} \Delta, \Gamma \vdash \text{case}(\underline{e}) x \cdot \underline{e}_1 y \cdot \underline{e}_2 : \\ \text{Pure } \underline{A}_2 ((\text{case}(e^*) x \cdot e_1^* y \cdot e_2^*) s_\Gamma) \end{aligned}$$

Which is our goal.

- (7) (ST-RET)

We have $\Delta \mid \Gamma \vdash e : A!n$. We need to show:

$$\Delta, \Gamma \vdash \text{return}_\tau e : \text{Pure } \underline{A} (\text{return}_\tau e^* s_\Gamma)$$

i.e.

$$\Delta, \Gamma \vdash \text{Pure.return } \underline{A} \underline{e} : \text{Pure } \underline{A} (\lambda p : A^* \rightarrow \text{Type}_0. (e^* s_\Gamma))$$

This is a trivial consequence of Theorem 14 by using the T-Ret rule of EMF^* , and the fact that $\underline{A} =_{\text{def}} A^*$.

- (8) (ST-BIND)

Suppose we have $\Delta \mid \Gamma \vdash e_1 : A! \tau$, and $\Delta \mid \Gamma, x : A \vdash e_2 : A'! \tau$, and so $\Delta \mid \Gamma \vdash \text{bind}_\tau e_1 \text{ to } x : A \text{ in } e_2 : A'! \tau$. We have to show:

$$\Delta, \Gamma \vdash \text{bind}_\tau e_1 \text{ to } x : A \text{ in } e_2 : \text{Pure } \underline{A}' ((\text{bind}_\tau e_1 \text{ to } x : A \text{ in } e_2)^* s_\Gamma)$$

that is:

$$\begin{aligned} \Delta, \Gamma \vdash \text{Pure.bind } \underline{A} \underline{A}' (e_1^* s_\Gamma) \underline{e}_2 (\lambda x : A^*. e_2^* s_\Gamma) (\lambda x : \underline{A}. \underline{e}_2) : \\ \text{Pure } \underline{A}' (\lambda p : A^* \rightarrow Ty. e_1^* (\lambda x : A^*. e_2^* p)) \end{aligned}$$

By our IHs we have:

$$\begin{aligned} \Delta, \Gamma \vdash \underline{e}_1 : \text{Pure } \underline{A} (e_1^* s_\Gamma) \\ \Delta, \Gamma, x : \underline{A} \vdash \underline{e}_2 : \text{Pure } \underline{A}' (e_2^* s_\Gamma) \end{aligned}$$

So we get:

$$\Delta, \Gamma \vdash \lambda x : \underline{A}. \underline{e}_2 : x : \underline{A} \rightarrow \text{Pure } \underline{A}' (e_2^* s_\Gamma)$$

By the T-Bind rule we can conclude:

$$\begin{aligned} \Delta, \Gamma \vdash \text{Pure.bind } \underline{A} \underline{A}' (e_1^* s_\Gamma) \underline{e}_2 (\lambda x : A^*. e_2^* s_\Gamma) (\lambda x : \underline{A}. \underline{e}_2) : \\ \text{Pure } \underline{A}' (\lambda p : \underline{A}' \rightarrow Ty. e_1^* (\lambda x : \underline{A}. \underline{e}_2^* p)) \end{aligned}$$

Since $\underline{A} = A^*$ and $\underline{A}' = A'^*$ this is exactly our goal.

Note: in this proof, we didn't use any specific fact about Pure , except the relation between monad operations and their WPs, so this is all generalizable to another target monad that's already defined and satisfies the base conditions for return and bind.

A.6 Equality preservation

We want to show that any source monad will give rise to specification-level monads in the target. This will be a consequence on the fact that equality is preserved by the $\star$ -translation. From equality preservation we also get the property that lifts are monad morphisms without further effort.

First we define equality for the source language. It is basically standard $\beta\eta$ -equivalence adding the monad laws for the base monad T . We keep the type and effect of each equality. It is an invariant that if we can derive an equality, both sides are well-typed at the specified type and effect.

The definition of the equality judgment is in Figure 8. Besides those rules, there is a congruence rule for every source construct, as expected.

We then prove that:

$$\frac{\Delta \mid \Gamma \vdash e_1 = e_2 : H! \varepsilon}{\Delta, \Gamma \models e_1^* s_\Gamma = e_2^* s_\Gamma}$$

Where by $\models$ it is meant the validity judgment of EMF^* .

Theorem 17 (Preservation of equality by CPS). *If $\Delta \mid \Gamma \vdash e_1 = e_2 : H! \varepsilon$ for any $\Delta, \Gamma, e_1, e_2, H, \varepsilon$, then one has $\Delta, \Gamma \models e_1^* s_\Gamma = e_2^* s_\Gamma$.*

Proof. By induction on the equality derivation. Most of the cases are trivial, since EMF^* has very similar rules for equality. The interesting cases are the monadic equalities, which we show here:

- (1) (EQ-M1)

We concluded

$$\Delta \mid \Gamma \vdash \text{bind}_\tau m \text{ to } x \text{ in } (\text{return}_\tau x) = m : A! \tau$$

Thus we need to show that

$$\Delta, \Gamma \models (\text{bind}_\tau m \text{ to } x \text{ in } (\text{return}_\tau x))^* s_\Gamma = m^* s_\Gamma$$

That is:

$$\Delta, \Gamma \models (\lambda p. (m^* s_\Gamma)) (\lambda x. (\lambda p'. p' x) p)) = m^* s_\Gamma$$

This is trivially provable by $\beta\eta$ -reduction.

- (2) (EQ-M2)

We concluded

$$\Delta \mid \Gamma \vdash \text{bind}_\tau (\text{return}_\tau e) \text{ to } x \text{ in } f x = f e : A'! \tau$$

thus we need to show that

$$\Delta, \Gamma \models (\text{bind}_\tau (\text{return}_\tau e) \text{ to } x \text{ in } f x)^* s_\Gamma = (f e)^* s_\Gamma$$

That is:

$$\Delta, \Gamma \models (\lambda p. (\lambda p'. p' e^*) (\lambda x. f^* x p)) s_\Gamma = (f^* s_\Gamma) (e^* s_\Gamma)$$

Note that since $x \notin FV(f) \implies x \notin FV(f^* s_\Gamma)$, this is easily shown by $\beta\eta$ -reduction as well.

- (3) (EQ-M3)

We concluded

$$\begin{aligned} \Delta \mid \Gamma \vdash \text{bind}_\tau (\text{bind}_\tau m \text{ to } x \text{ in } e_1) \text{ to } y \text{ in } e_2 \\ = \text{bind}_\tau m \text{ to } x \text{ in } (\text{bind}_\tau e_1 \text{ to } y \text{ in } e_2) : A''! \tau \end{aligned}$$

thus we need to show that

$$\begin{aligned} \Delta, \Gamma \models (\text{bind}_\tau (\text{bind}_\tau m \text{ to } x \text{ in } e_1) \text{ to } y \text{ in } e_2)^* s_\Gamma \\ = (\text{bind}_\tau m \text{ to } x \text{ in } (\text{bind}_\tau e_1 \text{ to } y \text{ in } e_2))^* s_\Gamma \end{aligned}$$

$$\begin{array}{c}
\text{EQ-BETA} \\
\frac{\Delta \mid \Gamma, x : H \vdash e_1 : H' ! \varepsilon \quad \Delta \mid \Gamma \vdash e_2 : H ! n}{\Delta \mid \Gamma \vdash (\lambda x : H. e_1) e_2 = e_1[e_2/x] : H' ! \varepsilon} \\
\\
\text{EQ-APP} \\
\frac{\Delta \mid \Gamma \vdash e_1 = e'_1 : H \xrightarrow{\varepsilon} H' ! n \quad \Delta \mid \Gamma \vdash e_2 = e'_2 : H ! n}{\Delta \mid \Gamma \vdash e_1 e_2 = e'_1 e'_2 : H' ! \varepsilon} \\
\\
\text{EQ-REFL} \quad \text{EQ-SYMM} \\
\frac{\Delta \mid \Gamma \vdash e : H ! \varepsilon}{\Delta \mid \Gamma \vdash e = e : H ! \varepsilon} \quad \frac{\Delta \mid \Gamma \vdash e_1 = e_2 : H ! \varepsilon}{\Delta \mid \Gamma \vdash e_2 = e_1 : H ! \varepsilon} \\
\\
\text{EQ-PAIR} \\
\frac{\Delta \mid \Gamma \vdash e : H \times H' ! n}{\Delta \mid \Gamma \vdash (\text{fst}(e), \text{snd}(e)) = e : H \times H' ! n} \\
\\
\text{EQ-M1} \\
\frac{\Delta \mid \Gamma \vdash m : A ! \tau}{\Delta \mid \Gamma \vdash \text{bind}_\tau m \text{ to } x \text{ in } (\text{return}_\tau x) = m : A ! \tau} \\
\\
\text{EQ-M2} \\
\frac{\Delta \mid \Gamma \vdash e : A ! n \quad \Delta \mid \Gamma \vdash f : A \xrightarrow{\varepsilon} A' ! n \quad x \notin FV(f)}{\Delta \mid \Gamma \vdash \text{bind}_\tau (\text{return}_\tau e) \text{ to } x \text{ in } f x = f e : A' ! \tau} \\
\\
\text{EQ-M3} \\
\frac{\Delta \mid \Gamma \vdash m : A ! \tau \quad \Delta \mid \Gamma, x : A \vdash e_1 : A' ! \tau \quad \Delta \mid \Gamma, y : A' \vdash e_2 : A'' ! \tau \quad x \notin FV(e_2)}{\Delta \mid \Gamma \vdash \text{bind}_\tau (\text{bind}_\tau m \text{ to } x \text{ in } e_1) \text{ to } y \text{ in } e_2 = \text{bind}_\tau m \text{ to } x \text{ in } (\text{bind}_\tau e_1 \text{ to } y \text{ in } e_2) : A'' ! \tau} \\
\\
\text{EQ-ETA} \\
\frac{\Delta \mid \Gamma \vdash e : H \xrightarrow{\varepsilon} H' ! n \quad x \notin FV(e)}{\Delta \mid \Gamma \vdash (\lambda x : H. e x) = e : H \xrightarrow{\varepsilon} H' ! n} \\
\\
\text{EQ-ABS} \\
\frac{\Delta \mid \Gamma, x : H \vdash e = e' : H' ! \varepsilon}{\Delta \mid \Gamma \vdash (\lambda x : H. e) = (\lambda x : H. e') : H \xrightarrow{\varepsilon} H' ! n} \\
\\
\text{EQ-TRANS} \\
\frac{\Delta \mid \Gamma \vdash e_1 = e_2 : H ! \varepsilon \quad \Delta \mid \Gamma \vdash e_2 = e_3 : H ! \varepsilon}{\Delta \mid \Gamma \vdash e_1 = e_3 : H ! \varepsilon} \\
\\
\text{EQ-CASE} \\
\frac{\Delta \mid \Gamma \vdash e : A + A' ! n}{\Delta \mid \Gamma \vdash \text{case } e \text{ inl } x. \text{inl}(x); \text{inr } x. \text{inr}(x) = e : A + A' ! n}
\end{array}$$

Figure 8. Equality rules for DM

That is:

$$\begin{aligned}
\Delta, \Gamma \models (\lambda p. (\lambda p'. (m^* s_\Gamma) (\lambda x. (e_1^* s_\Gamma) p')) (\lambda y. (e_2^* s_\Gamma) p)) \\
= (\lambda p. (m^* s_\Gamma) (\lambda x. (\lambda p'. (e_1^* s_\Gamma) (\lambda y. (e_2^* s_\Gamma) p')) p))
\end{aligned}$$

Note that since $x \notin FV(e_2) \implies x \notin FV(e_2^*)$, this is also easily shown by $\beta\eta$ -reduction.

The proof above is easy, and that should not be surprising, as we are translating our abstract monadic operations into a concrete monad (continuations), thus our source equalities should be trivially satisfied after translation.

□

A.7 Monotonicity

We're interested in the monotonicity of WPs. Firstly, we need a higher-order definition for this property. Throughout this section we mostly ignore the image of our $\star$ -translation and work with a larger subset of EMF * language. This gives us a stronger result than strictly necessary.

The types where the translation is defined are those non-dependent and monad-free (meaning every arrow in them is a Tot-arrow). No occurrence of Pure is allowed. The types of specifications are always of this shape, so this is not a limitation.

For non-empty environments the theorem states:

Theorem 18 (Monotonicity of $\star$ -translation— environments). *For any Δ, Γ, e, H, A one has:*

1. $\Delta \mid \Gamma \vdash e : H ! n \implies \Delta, \Gamma^{12} \models e^{*1} \lesssim_{H^*} e^{*2}$
2. $\Delta \mid \Gamma \vdash e : A ! \tau \implies \Delta, \Gamma^{12} \models e^{*1} \lesssim_{(A^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} e^{*2}$

Where if $\Gamma = x_1 : t_1, \dots, x_n : t_n$ we define $\Gamma^{12} = x_1^1 : t_1^1, x_1^2 : t_1^2, [x_1 \lesssim_{t_1^1} x_2], \dots$ essentially duplicating each variable and introducing a strengthening hypothesis between them ($[\phi]$ is notation for $h : \phi$, where h does not appear free in the RHS). We then define the “1” substitution as $[x_1^1/x_1, \dots, x_n^1/x_n]$ and similarly for “2”. This trivially implies both previous monotonicity theorems.

Proof. We prove these two propositions by induction on the typing derivation for e . Throughout the proof Δ plays no special role, so we just drop it from the reasoning, keeping in mind that it has to be there for having well-formed types (but nothing else).

Note that during the proof we treat $\lesssim_X$ abstractly, so any instantiation with a proper type (not necessarily those where $\lesssim$ reduces to equality) would be OK.

Throughout this proof we sometimes skip the subindices for $\lesssim$ in favor of compactness. Hopefully, they should be clear from the context.

(1) (ST-VAR)

We need to show $x_i^1 \lesssim_{t_i^*} x_i^2$. This is trivial from the context and by using the (V-ASSUME) rule.

(2) (ST-CONST)

The constants only deal with base types, so all inductive hypotheses for the arguments reduce to an equality, as does our goal. Our goal is then trivially provable by applications of (V-EQP).

(3) (ST-ABS)

Say we concluded $\Gamma, x : t \vdash e : s ! \varepsilon$. As IH we have:

$$\Gamma^{12}, x^1 : t^*, x^2 : t^*, [x^1 \lesssim_{t^*} x^2] \models e^{*1} [x^1/x] \lesssim_{s'} e^{*2} [x^2/x]$$

Where s' is either s^* or $(s^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$ depending on ε . The proof is independent of this. What we need to prove is:

$$\Gamma^{12} \models (\lambda x : t^*. e^{*1}) \lesssim_{t^* \rightarrow s'} (\lambda x : t^*. e^{*2})$$

Which by definition is:

$$\Gamma^{12} \models \forall x^1, x^2 : t^*, x^1 \lesssim_{t^*} x^2 \implies (\lambda x : t^*. e^{*1}) x^1 \lesssim_{s'} (\lambda x : t^*. e^{*2}) x^2$$

By reduction ((V-EQRED) + (V-EQ*)), this is equivalent to:

$$\Gamma^{12} \models \forall x^1, x^2 : t^*, x^1 \lesssim_{t^*} x^2 \implies e^{*1} [x^1/x] \lesssim_{s'} e^{*2} [x^2/x]$$

Which we can conclude from our IH and three applications on (V- $\forall$ 1).

(4) (ST-APP)

Say $\Gamma \vdash f : t \xrightarrow{\varepsilon} s ! n$ and $\Gamma \vdash e : a ! n$. As inductive hypothesis we get:

$$\Gamma^{12} \models f^{*1} \lesssim_{t^* \rightarrow s'} f^{*2} \quad \Gamma^{12} \models e^{*1} \lesssim_{t^*} e^{*2}$$

Where s' is either s^* or $(s^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$ depending on ε . Again, the proof is independent of this. Expanding the definition of $\lesssim$ on the left we get:

$$\Gamma^{12} \models \forall x^1, x^2 : t^*, x^1 \lesssim_{t^*} x^2 \implies f^{*1} x^1 \lesssim_{s'} f^{*2} x^2$$

We instantiate (using (V- $\forall$ E)) x^1, x^2 with e^{*1}, e^{*2} , and apply (V-MP) with our other IH to get:

$$\Gamma^{12} \models f^{*1} e^{*1} \lesssim_{s'} f^{*2} e^{*2}$$

Which is exactly our goal in any $(\xrightarrow{n}, \xrightarrow{\tau})$ case.

(5) (ST-RET)

Say $\Gamma \vdash \text{return}_\tau e : t ! \tau$. Our IH gives us:

$$\Gamma^{12} \models e^{*1} \lesssim_{t^*} e^{*2}$$

And we need to show that:

$$\Gamma^{12} \models (\lambda p. p e^{*1}) \lesssim_{(t^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} (\lambda p. p e^{*2})$$

That is:

$$\Gamma^{12} \models \forall p^1, p^2, p^1 \lesssim p^2 \implies (\lambda p. p e^{*1}) p^1 \lesssim_{\text{Type}_0} (\lambda p. p e^{*2}) p^2$$

By reduction:

$$\Gamma^{12} \models \forall p^1, p^2, p^1 \lesssim p^2 \implies p^1 e^{*1} \lesssim_{\text{Type}_0} p^2 e^{*2}$$

Which is trivially provable by preservation of $\lesssim$ by application and the IH for e .

(6) (ST-BIND)

Say $\Gamma \vdash m : a ! \tau$ and $\Gamma, x : a \vdash e : b ! \tau$, so we get $\Gamma \vdash \text{bind}_\tau m \text{ to } x \text{ in } e : b ! \tau$. Our IHs are:

$$\Gamma^{12} \models m^{*1} \lesssim_{(a \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} m^{*2} \quad \Gamma^{12}, x^1 : a^*, x^2 : a^*, [x^1 \lesssim_{a^*} x^2] \models e^{*1} [x^1/x] \lesssim_{(b \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} e^{*2} [x^2/x]$$

We need to show that:

$$\Gamma^{12} \models (\lambda p. m^{*1} (\lambda x. e^{*1} p)) \lesssim_{(b^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0} (\lambda p. m^{*2} (\lambda x. e^{*2} p))$$

Which can be simplified to:

$$\Gamma^{12}, p^1, p^2, [p^1 \lesssim p^2] \models m^{*1} (\lambda x. e^{*1} p^1) \lesssim_{\text{Type}_0} m^{*2} (\lambda x. e^{*2} p^2)$$

Since $m^{*1} \lesssim m^{*2}$ by the IH, this can be concluded by:

$$\Gamma^{12}, p^1, p^2, [p^1 \lesssim p^2] \models \lambda x. e^{*1} p^1 \lesssim_{a^* \rightarrow \text{Type}_0} \lambda x. e^{*2} p^2$$

Which can be simplified to:

$$\Gamma^{12}, p^1, p^2, [p^1 \lesssim p^2], x^1 : a^*, x^2 : a^*, [x^1 \lesssim_{a^*} x^2] \models e^{*1} [x^1/x] p^1 \lesssim_{\text{Type}_0} e^{*2} [x^2/x] p^2$$

Weakening the IH for e we know $e^{*1} [x^1/x] \lesssim e^{*2} [x^2/x]$ in this context. Since we also have $p^1 \lesssim p^2$ and $\lesssim$ is preserved by application we have our goal.

(7) (ST-PAIR), (ST-FST), (ST-INL)

All trivial from IHs.

(8) (ST-CASE)

By case analysis on the IH for the sum type, and reduction.

□

Having this proof implies that any well-typed term will be given a monotonic specification. And, as a consequence, functions preserve monotonicity.

A.8 Conjunctivity

The definition of conjunctivity on EMF* predicate types was given previously. The full theorem which we prove is this:

Theorem 19 (Conjunctivity of $\star$ -translation— environments). *For any Δ, Γ, e, H, A one has:*

1. $\Delta \mid \Gamma \vdash e : C ! n \implies \Delta, \Gamma_{\mathbb{C}} \models \mathbb{C}_{C^*}(e^*)$
2. $\Delta \mid \Gamma \vdash e : A ! \tau \implies \Delta, \Gamma_{\mathbb{C}} \models \mathbb{C}_{(A^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0}(e^*)$

Where when $\Gamma = x_1 : t_1, \dots$, we define $\Gamma_{\mathbb{C}} = x_1 : t_1^*, [\mathbb{C}_{t_1^*}(x_1)], \dots$. This trivially implies the previously stated theorem by taking $\Gamma = \cdot$.

Proof. By induction on the typing derivations. Once again, Δ does not play a big role and we omit it.

(1) (ST-VAR)

Trivial from context, for any type.

(2) (ST-CONST)

Does not apply as no constant gives a type $C ! n$ nor $A ! \tau$

(3) (ST-ABS)

Say we concluded $\Gamma, x : t \vdash e : s ! \varepsilon$ (where that might be $C ! n$ or $A ! \tau$, we treat both cases uniformly). From the IH we get

$$\Gamma_{\mathbb{C}}, x : t^*, [\mathbb{C}_{t^*}(x)] \models \mathbb{C}_{s^*}(e)$$

Where s' is s^* or $(s^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0$ according to (s, ε) . By applying (V- $\forall$ 1) twice we get:

$$\Gamma_{\mathbb{C}} \models \forall x : t^*. \mathbb{C}_{t^*}(x) \implies \mathbb{C}_{s'}(e)$$

Which is the same, by reduction, as:

$$\Gamma_{\mathbb{C}} \models \forall x : t^*. \mathbb{C}_{t^*}(x) \implies \mathbb{C}_{s'}((\lambda x. e x) x)$$

Thus by definition of $\mathbb{C}$:

$$\Gamma_{\mathbb{C}} \models \mathbb{C}_{t \rightarrow s'}(\lambda x. e x)$$

As required for both cases.

(4) (ST-APP)

Trivial by the preservation of $\mathbb{C}$ by application, in both cases (applies (V-MP)).

(5) (ST-RET)

Say we concluded $\Gamma \vdash \text{return}_\tau e : A ! \tau$. Our goal is then:

$$\Gamma_{\mathbb{C}} \models \mathbb{C}_{(A^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0}(\lambda p. p e^*)$$

Which is:

$$\Gamma_{\mathbb{C}} \models \forall p_1, p_2. (\lambda p. p e^*) p_1 \wedge (\lambda p. p e^*) p_2 = (\lambda p. p e^*)(\lambda x. p_1 x \wedge p_2 x)$$

By reduction that's equivalent to:

$$\Gamma_{\mathbb{C}} \models \forall p_1, p_2. p_1 e^* \wedge p_2 e^* = p_1 e^* \wedge p_2 e^*$$

Which is trivially true (without use of any IH) by (V-REFL).

(6) (ST-BIND)

Say we concluded $\Gamma \vdash \text{bind}_\tau e_1 \text{ to } x \text{ in } e_2 : A' ! \tau$, where $e_1 : A ! \tau$. Our IHs are:

$$\Gamma_{\mathbb{C}} \models \mathbb{C}_{(A^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0}(e_1^*) \quad \Gamma_{\mathbb{C}}, x : A^*, [\mathbb{C}_{A^*}(x)] \models \mathbb{C}_{(A'^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0}(e_2^*)$$

We need to show:

$$\Gamma_{\mathbb{C}} \models \mathbb{C}_{(A'^* \rightarrow \text{Type}_0) \rightarrow \text{Type}_0}(\lambda p. e_1^*(\lambda x. e_2^* p))$$

Expanding the definition, this is:

$$\Gamma_{\mathbb{C}} \models \forall p_1, p_2. (\lambda p. e_1^*(\lambda x. e_2^* p)) p_1 \wedge (\lambda p. e_1^*(\lambda x. e_2^* p)) p_2 = (\lambda p. e_1^*(\lambda x. e_2^* p))(\lambda x. p_1 x \wedge p_2 x)$$

By reduction, this is equivalent to:

$$\Gamma_{\mathbb{C}} \models \forall p_1, p_2. e_1^*(\lambda x. e_2^* p_1) \wedge e_1^*(\lambda x. e_2^* p_2) = e_1^*(\lambda x. e_2^*(\lambda x. p_1 x \wedge p_2 x))$$

By the IH for e_2 we know $\forall x. e_2^* (\lambda x. p_1 x \wedge p_2 x) = e_2^* p_1 \wedge e_2^* p_2$. By reduction and (V-EXT) this means $(\lambda x. e_2^* (\lambda x. p_1 x \wedge p_2 x)) = (\lambda x. e_2^* p_1 \wedge e_2^* p_2)$. Thus we replace on the RHS (via (V-SUBST)) and get:

$$\Gamma_{\mathbb{C}} \models \forall p_1, p_2. e_1^* (\lambda x. e_2^* p_1) \wedge e_1^* (\lambda x. e_2^* p_2) = e_1^* (\lambda x. e_2^* p_1 \wedge e_2^* p_2)$$

By some η -expansion and the IH for e_1 we can turn this to:

$$\Gamma_{\mathbb{C}} \models \forall p_1, p_2. e_1^* (\lambda x. e_2^* p_1) \wedge e_1^* (\lambda x. e_2^* p_2) = e_1^* (\lambda x. e_2^* p_1) \wedge e_1^* (\lambda x. e_2^* p_2)$$

Which is trivially provable by (V-REFL).

- (7) (ST-PAIR), (ST-FST)
All trivial by IHs.
- (8) (ST-INL)
Does not apply for the cases we consider.
- (9) (ST-CASE)
Trivial by (V-SUMIND) and the IHs.

□

Thus, any term obtained by the $\star$ -translation (return, bind, actions, lifts, ...) will be conjunctive in this sense, which means they also preserve the property through application.

With a completely analogous definition and proof we get the expected result of conjunctivity over (non-empty) universal quantification. The non-empty requirement is not actually stressed during that proof, but it's the wanted result as WPs (which can be taken as arguments) might not distribute over empty universals.

References

- D. Ahman and T. Uustalu. Update monads: Cointerpreting directed containers. *TYPES*, 2013.
- D. Ahman, N. Ghani, and G. D. Plotkin. Dependent types and fibred computational effects. *FOSSACS*, 2016.
- R. Atkey. Parameterised notions of computation. *Journal of Functional Programming*, 19:335–376, 2009.
- P. Audebaud and E. Zucca. Deriving proof rules from continuation semantics. *Formal Asp. Comput.*, 11(4):426–447, 1999.
- G. Barthe, C. Fournet, B. Grégoire, P.-Y. Strub, N. Swamy, and S. Zanella-Béguélin. Probabilistic relational verification for cryptographic implementations. *POPL*, 2014.
- N. Benton. Simple relational correctness proofs for static analyses and program transformations. *POPL*, 2004.
- N. Benton and A. Kennedy. Exceptional syntax. *J. Funct. Program.*, 11(4):395–410, 2001.
- N. Benton, J. Hughes, and E. Moggi. *Monads and Effects*, pages 42–122. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.
- E. Brady. Programming and reasoning with algebraic effects and dependent types. *ICFP*, 2013.
- C. Casinghino, V. Sjöberg, and S. Weirich. Combining proofs and programs in a dependently typed language. *POPL*, 2014.
- A. Charguéraud. Characteristic formulae for the verification of imperative programs. *ICFP*, 2011.
- A. Chlipala, G. Malecha, G. Morrisett, A. Shinnar, and R. Wisnesky. Effective interactive proofs for higher-order imperative programs. *ICFP*, 2009.
- T. Coquand and G. Huet. The calculus of constructions. *Information and Computation*, 76(2):95 – 120, 1988.
- L. M. de Moura and N. Björner. Z3: an efficient SMT solver. *TACAS*, 2008.
- E. W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM*, 18(8):453–457, 1975.
- E. W. Dijkstra. *A Discipline of Programming*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1st edition, 1997.
- A. Filinski. Representing monads. *POPL*, 1994.
- A. Filinski. Representing layered monads. *POPL*, 1999.
- A. Filinski. Monads in action. *POPL*, 2010.
- J.-C. Filliâtre and A. Paskevich. Why3 — where programs meet provers. *ESOP*, 2013.
- C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. *PLDI*, 1993.
- B. Jacobs. Dijkstra and Hoare monads in monadic computation. *Theor. Comput. Sci.*, 604:30–45, 2015.
- K. Jensen. Connection between Dijkstra’s predicate-transformers and denotational continuation-semantics. DAIMI Report Series 7.86, 1978.
- K. Keimel. Healthiness conditions for predicate transformers. *Electr. Notes Theor. Comput. Sci.*, 319:255–270, 2015.
- K. Keimel and G. Plotkin. Mixed powerdomains for probability and nondeterminism. submitted to LMCS, 2016.
- K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. *LPAR*, 2010.
- E. Moggi. Computational lambda-calculus and monads. *LICS*, 1989.
- A. Nanevski, J. G. Morrisett, and L. Birkedal. Hoare type theory, polymorphism and separation. *JFP*, 18(5-6):865–911, 2008.
- A. Nogin. *Quotient Types: A Modular Approach*, pages 263–280. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.
- C. Paulin-Mohring. Introduction to the Calculus of Inductive Constructions. In B. W. Paleo and D. Delahaye, editors, *All about Proofs, Proofs for All*, volume 55 of *Studies in Logic (Mathematical logic and foundations)*. College Publications, 2015.
- G. Plotkin and M. Pretnar. *Handlers of Algebraic Effects*, pages 80–94. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE J.Sel. A. Commun.*, 21(1):5–19, 2006.
- D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières. Flexible dynamic information flow control in haskell. *SIGPLAN Not.*, 46(12):95–106, 2011.
- N. Swamy, N. Guts, D. Leijen, and M. Hicks. Lightweight monadic programming in ML. *ICFP*, 2011.
- N. Swamy, J. Weinberger, C. Schlesinger, J. Chen, and B. Livshits. Verifying higher-order programs with the Dijkstra monad. *PLDI*, 2013.
- N. Swamy, C. Hrițcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, P.-Y. Strub, M. Kohlweiss, J.-K. Zinzindohoue, and S. Z. Béguélin. Dependent types and multi-monadic effects in F*. *POPL*, 2016.
- P. Wadler. Comprehending monads. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming*. 1990.
- P. Wadler. The essence of functional programming. *POPL*, 1992.
- P. Wadler. Monads and composable continuations. *Lisp Symb. Comput.*, 7(1):39–56, 1994.