



Explicit Effect Subtyping

Amr Hany Saleh¹(✉), Georgios Karachalias¹, Matija Pretnar²,
and Tom Schrijvers¹

¹ Department of Computer Science, KU Leuven, Leuven, Belgium
ah.saleh@cs.kuleuven.be

² Faculty of Mathematics and Physics, University of Ljubljana, Ljubljana, Slovenia

Abstract. As popularity of algebraic effects and handlers increases, so does a demand for their efficient execution. Eff, an ML-like language with native support for handlers, has a subtyping-based effect system on which an effect-aware optimizing compiler could be built. Unfortunately, in our experience, implementing optimizations for Eff is overly error-prone because its core language is implicitly-typed, making code transformations very fragile.

To remedy this, we present an explicitly-typed polymorphic core calculus for algebraic effect handlers with a subtyping-based type-and-effect system. It reifies appeals to subtyping in explicit casts with coercions that witness the subtyping proof, quickly exposing typing bugs in program transformations.

Our typing-directed elaboration comes with a constraint-based inference algorithm that turns an implicitly-typed Eff-like language into our calculus. Moreover, all coercions and effect information can be erased in a straightforward way, demonstrating that coercions have no computational content.

1 Introduction

Algebraic effect handlers [17, 18] are quickly maturing from a theoretical model to a practical language feature for user-defined computational effects. Yet, in practice they still incur a significant performance overhead compared to native effects.

Our earlier efforts [22] to narrow this gap with an optimising compiler from Eff [2] to OCaml showed promising results, in some cases reaching even the performance of hand-tuned code, but were very fragile and have been postponed until a more robust solution is found. We believe the main reason behind this fragility is the complexity of subtyping in combination with the implicit typing of Eff’s core language, further aggravated by the “garbage collection” of subtyping constraints (see Sect. 7).¹

¹ For other issues stemming from the same combination see issues #11 and #16 at <https://github.com/matijapretnar/eff/issues/>.

For efficient compilation, one must avoid the poisoning problem [26], where unification forces a pure computation to take the less precise impure type of the context (e.g. a pure and an impure branch of a conditional both receive the same impure type). Since this rules out existing (and likely simpler) effect systems for handlers based on row-polymorphism [8, 12, 14], we propose a polymorphic explicitly-typed calculus based on subtyping. More specifically, our contributions are as follows:

- First, in Sect. 3 we present IMPEFF, a polymorphic implicitly-typed calculus for algebraic effects and handlers with a subtyping-based type-and-effect system. IMPEFF is essentially a (desugared) source language as it appears in the compiler frontend of a language like Eff.
- Next, Sect. 4 presents EXEFF, the core calculus, which combines explicit System F-style polymorphism with explicit coercions for subtyping in the style of Breazu-Tannen et al. [3]. This calculus comes with a type-and-effect system, a small-step operational semantics and a proof of type-safety.
- Section 5 specifies the typing-directed elaboration of IMPEFF into EXEFF and presents a type inference algorithm for IMPEFF that produces the elaborated EXEFF term as a by-product. It also establishes that the elaboration preserves typing, and that the algorithm is sound with respect to the specification and yields principal types.
- Finally, Sect. 6 defines SKELEFF, which is a variant of EXEFF without effect information or coercions. SKELEFF is also representative of Multicore Ocaml’s support for algebraic effects and handlers [6], which is a possible compilation target of Eff. By showing that the erasure from EXEFF to SKELEFF preserves semantics, we establish that EXEFF’s coercions are computationally irrelevant and that, despite the existence of multiple proofs for the same subtyping, there is no coherence problem. To enable erasure, EXEFF annotates its types with (*type*) *skeletons*, which capture the erased counterpart and are, to our knowledge, a novel contribution.
- Our paper comes with two software artefacts: an ongoing implementation² of a compiler from Eff to OCaml with EXEFF at its core, and an Abella mechanisation³ of Theorems 1, 2, 6, and 7. Remaining theorems all concern the inference algorithm, and their proofs closely follow [20].

The full version of this paper includes an appendix with omitted figures and can be found at <http://www.cs.kuleuven.be/publicaties/rapporten/cw/CW711.abs.html>.

2 Overview

This section presents an informal overview of the EXEFF calculus, and the main issues with elaborating to and erasing from it.

² <https://github.com/matijapretnar/eff/tree/explicit-effect-subtyping>.

³ <https://github.com/matijapretnar/proofs/tree/master/explicit-effect-subtyping>.

2.1 Algebraic Effect Handlers

The main premise of algebraic effects is that impure behaviour arises from a set of *operations* such as `Get` and `Set` for mutable store, `Read` and `Print` for interactive input and output, or `Raise` for exceptions [17]. This allows generalizing exception handlers to other effects, to express backtracking, co-operative multithreading and other examples in a natural way [2, 18].

Assume operations $\text{Tick} : \text{Unit} \rightarrow \text{Unit}$ and $\text{Tock} : \text{Unit} \rightarrow \text{Unit}$ that take a unit value as a parameter and yield a unit value as a result. Unlike special built-in operations, these operations have no intrinsic effectful behaviour, though we can give one through handlers. For example, the handler $\{\text{Tick } x \ k \mapsto (\text{Print } \text{“tick”}; k \ \text{unit}), \text{Tock } x \ k \mapsto \text{Print } \text{“tock”}\}$ replaces all calls of `Tick` by printing out “tick” and similarly for `Tock`. But there is one significant difference between the two cases. Unlike exceptions, which always abort the evaluation, operations have a continuation waiting for their result. It is this continuation that the handler captures in the variable k and potentially uses in the handling clause. In the clause for `Tick`, the continuation is resumed by passing it the expected unit value, whereas in the clause for `Tock`, the operation is discarded. Thus, if we handle a computation emitting the two operations, it will print out “tick” until a first “tock” is printed, after which the evaluation stops.

2.2 Elaborating Subtyping

Consider the computation $\text{do } x \leftarrow \text{Tick } \text{unit}; f \ x$ and assume that f has the function type $\text{Unit} \rightarrow \text{Unit} ! \{\text{Tock}\}$, taking unit values to unit values and perhaps calling `Tock` operations in the process. The whole computation then has the type $\text{Unit} ! \{\text{Tick}, \text{Tock}\}$ as it returns the unit value and may call `Tick` and `Tock`.

The above typing implicitly appeals to subtyping in several places. For instance, $\text{Tick } \text{unit}$ has type $\text{Unit} ! \{\text{Tick}\}$ and $f \ x$ type $\text{Unit} ! \{\text{Tock}\}$. Yet, because they are sequenced with `do`, the type system expects they have the same set of effects. The discrepancies are implicitly reconciled by the subtyping which admits both $\{\text{Tick}\} \leq \{\text{Tick}, \text{Tock}\}$ and $\{\text{Tock}\} \leq \{\text{Tick}, \text{Tock}\}$.

We elaborate the `IMPEFF` term into the explicitly-typed core language `EXEFF` to make those appeals to subtyping explicit by means of casts with coercions:

$$\text{do } x \leftarrow ((\text{Tick } \text{unit}) \triangleright \gamma_1); (f \ x) \triangleright \gamma_2$$

A coercion γ is a witness for a subtyping $A ! \Delta \leq A' ! \Delta'$ and can be used to cast a term c of type $A ! \Delta$ to a term $c \triangleright \gamma$ of type $A' ! \Delta'$. In the above term, γ_1 and γ_2 respectively witness $\text{Unit} ! \{\text{Tick}\} \leq \text{Unit} ! \{\text{Tick}, \text{Tock}\}$ and $\text{Unit} ! \{\text{Tock}\} \leq \text{Unit} ! \{\text{Tick}, \text{Tock}\}$.

2.3 Polymorphic Subtyping for Types and Effects

The above basic example only features monomorphic types and effects. Yet, our calculus also supports polymorphism, which makes it considerably more

expressive. For instance the type of f in `let f = (fun g ↦ g unit) in ...` is generalised to:

$$\forall \alpha, \alpha'. \forall \delta, \delta'. \alpha \leq \alpha' \Rightarrow \delta \leq \delta' \Rightarrow (\mathbf{Unit} \rightarrow \alpha ! \delta) \rightarrow \alpha' ! \delta'$$

This polymorphic type scheme follows the qualified types convention [9] where the type $(\mathbf{Unit} \rightarrow \alpha ! \delta) \rightarrow \alpha' ! \delta'$ is subjected to several qualifiers, in this case $\alpha \leq \alpha'$ and $\delta \leq \delta'$. The universal quantifiers on the outside bind the type variables α and α' , and the effect set variables δ and δ' .

The elaboration of f into EXEFF introduces explicit binders for both the quantifiers and the qualifiers, as well as the explicit casts where subtyping is used.

$$\Lambda \alpha. \Lambda \alpha'. \Lambda \delta. \Lambda \delta'. \Lambda (\omega : \alpha \leq \alpha'). \Lambda (\omega' : \delta \leq \delta'). \mathbf{fun} (g : \mathbf{Unit} \rightarrow \alpha ! \delta) \mapsto (g \mathbf{unit}) \triangleright (\omega ! \omega')$$

Here the binders for qualifiers introduce coercion variables ω between pure types and ω' between operation sets, which are then combined into a computation coercion $\omega ! \omega'$ and used for casting the function application $g \mathbf{unit}$ to the expected type.

Suppose that h has type $\mathbf{Unit} \rightarrow \mathbf{Unit} ! \{\mathbf{Tick}\}$ and $f h$ type $\mathbf{Unit} ! \{\mathbf{Tick}, \mathbf{Tock}\}$. In the EXEFF calculus the corresponding instantiation of f is made explicit through type and coercion applications

$$f \mathbf{Unit} \mathbf{Unit} \{\mathbf{Tick}\} \{\mathbf{Tick}, \mathbf{Tock}\} \gamma_1 \gamma_2 h$$

where γ_1 needs to be a witness for $\mathbf{Unit} \leq \mathbf{Unit}$ and γ_2 for $\{\mathbf{Tick}\} \leq \{\mathbf{Tick}, \mathbf{Tock}\}$.

2.4 Guaranteed Erasure with Skeletons

One of our main requirements for EXEFF is that its effect information and subtyping can be easily erased. The reason is twofold. Firstly, we want to show that neither plays a role in the runtime behaviour of EXEFF programs. Secondly and more importantly, we want to use a conventionally typed (System F-like) functional language as a backend for the Eff compiler.

At first, erasure of both effect information and subtyping seems easy: simply drop that information from types and terms. But by dropping the effect variables and subtyping constraints from the type of f , we get $\forall \alpha, \alpha'. (\mathbf{Unit} \rightarrow \alpha) \rightarrow \alpha'$ instead of the expected type $\forall \alpha. (\mathbf{Unit} \rightarrow \alpha) \rightarrow \alpha$. In our naive erasure attempt we have carelessly discarded the connection between α and α' . A more appropriate approach to erasure would be to unify the types in dropped subtyping constraints. However, unifying types may reduce the number of type variables when they become instantiated, so corresponding binders need to be dropped, greatly complicating the erasure procedure and its meta-theory.

Fortunately, there is an easier way by tagging all bound type variables with *skeletons*, which are barebone types without effect information. For example, the skeleton of a function type $A \rightarrow B ! \Delta$ is $\tau_1 \rightarrow \tau_2$, where τ_1 is the skeleton of

A and τ_2 the skeleton of B . In EXEFF every well-formed type has an associated skeleton, and any two types $A_1 \leq A_2$ share the same skeleton. In particular, binders for type variables are explicitly annotated with skeleton variables ς . For instance, the actual type of f is:

$$\forall \varsigma. \forall (\alpha : \varsigma), (\alpha' : \varsigma). \forall \delta, \delta'. \alpha \leq \alpha' \Rightarrow \delta \leq \delta' \Rightarrow (\mathbf{Unit} \rightarrow \alpha ! \delta) \rightarrow \alpha' ! \delta'$$

The skeleton quantifications and annotations also appear at the term-level:

$$\Lambda \varsigma. \Lambda (\alpha : \varsigma). \Lambda (\alpha' : \varsigma). \Lambda \delta. \Lambda \delta'. \Lambda (\omega : \alpha \leq \alpha'). \Lambda (\omega' : \delta \leq \delta'). \dots$$

Now erasure is really easy: we drop not only effect and subtyping-related term formers, but also type binders and application. We do retain skeleton binders and applications, which take over the role of (plain) types in the backend language. In terms, we replace types by their skeletons. For instance, for f we get:

$$\Lambda \varsigma. \mathbf{fun} (g : \mathbf{Unit} \rightarrow \varsigma) \mapsto g \mathbf{unit} \quad : \quad \forall \varsigma. (\mathbf{Unit} \rightarrow \varsigma) \rightarrow \varsigma$$

Terms

```

value  $v ::= x \mid \mathbf{unit} \mid \mathbf{fun} x \mapsto c \mid h$ 
handler  $h ::= \{\mathbf{return} x \mapsto c_r, \mathbf{Op}_1 x k \mapsto c_{\mathbf{Op}_1}, \dots, \mathbf{Op}_n x k \mapsto c_{\mathbf{Op}_n}\}$ 
computation  $c ::= \mathbf{return} v \mid \mathbf{Op} v (y.c) \mid \mathbf{do} x \leftarrow c_1; c_2$ 
               $\mid \mathbf{handle} c \mathbf{with} v \mid v_1 v_2 \mid \mathbf{let} x = v \mathbf{in} c$ 

```

Types & Constraints

```

skeleton  $\tau ::= \varsigma \mid \mathbf{Unit} \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \Rightarrow \tau_2$ 

value type  $A, B ::= \alpha \mid \mathbf{Unit} \mid A \rightarrow \underline{C} \mid \underline{C} \Rightarrow \underline{D}$ 
qualified type  $K ::= A \mid \pi \Rightarrow K$ 
polytype  $S ::= K \mid \forall \varsigma. S \mid \forall \alpha : \tau. S \mid \forall \delta. S$ 
computation type  $\underline{C}, \underline{D} ::= A ! \Delta$ 
dirt  $\Delta ::= \delta \mid \emptyset \mid \{\mathbf{Op}\} \cup \Delta$ 

simple constraint  $\pi ::= A_1 \leq A_2 \mid \Delta_1 \leq \Delta_2$ 
constraint  $\rho ::= \pi \mid \underline{C} \leq \underline{D}$ 

```

Fig. 1. IMPEFF Syntax

3 The ImpEff Language

This section presents IMPEFF, a basic functional calculus with support for algebraic effect handlers, which forms the core language of our optimising compiler. We describe the relevant concepts, but refer the reader to Pretnar's tutorial [21], which explains essentially the same calculus in more detail.

3.1 Syntax

Figure 1 presents the syntax of the source language. There are two main kinds of terms: (pure) values v and (dirty) computations c , which may call effectful operations. Handlers h are a subsidiary sort of values. We assume a given set of operations \mathbf{Op} , such as **Get** and **Put**. We abbreviate $\mathbf{Op}_1 x k \mapsto c_{\mathbf{Op}_1}, \dots, \mathbf{Op}_n x k \mapsto c_{\mathbf{Op}_n}$ as $[\mathbf{Op} x k \mapsto c_{\mathbf{Op}}]_{\mathbf{Op} \in \mathcal{O}}$, and write \mathcal{O} to denote the set $\{\mathbf{Op}_1, \dots, \mathbf{Op}_n\}$.

Similarly, we distinguish between two basic sorts of types: the value types A, B and the computation types $\underline{C}, \underline{D}$. There are four forms of value types: type variables α , function types $A \rightarrow \underline{C}$, handler types $\underline{C} \Rightarrow \underline{D}$ and the **Unit** type. Skeletons τ capture the shape of types, so, by design, their forms are identical. The computation type $A ! \Delta$ is assigned to a computation returning values of type A and potentially calling operations from the *dirt* set Δ . A dirt set contains zero or more operations \mathbf{Op} and is terminated either by an empty set or a dirt variable δ . Though we use **cons-list** syntax, the intended semantics of dirt sets Δ is that the order of operations \mathbf{Op} is irrelevant. Similarly to all HM-based systems, we discriminate between value types (or monotypes) A , qualified types K and polytypes (or type schemes) S . (Simple) subtyping constraints π denote inequalities between either value types or dirts. We also present the more general form of constraints ρ that includes inequalities between computation types (as we illustrate in Sect. 3.2 below, this allows for a single, uniform constraint entailment relation). Finally, polytypes consist of zero or more skeleton, type or dirt abstractions followed by a qualified type.

3.2 Typing

Figure 2 presents the typing rules for values and computations, along with a typing-directed elaboration into our target language EXEFF. In order to simplify the presentation, in this section we focus exclusively on typing. The parts of the rules that concern elaboration are highlighted in gray and are discussed in Sect. 5.

Values. Typing for values takes the form $\Gamma \vdash_v v : A \rightsquigarrow v'$, and, given a typing environment Γ , checks a value v against a value type A .

Rule **TMVAR** handles term variables. Given that x has type $(\forall \bar{\zeta}. \bar{\alpha} : \bar{\tau}. \forall \bar{\delta}. \bar{\pi} \Rightarrow A)$, we *appropriately* instantiate the skeleton $(\bar{\zeta})$, type $(\bar{\alpha})$, and dirt $(\bar{\delta})$ variables, and ensure that the instantiated wanted constraints $\sigma(\bar{\pi})$ are satisfied, via side condition $\bar{\Gamma} \vdash_{\sigma} \gamma : \sigma(\bar{\pi})$. Rule **TMCASTV** allows casting the type of a value v from A to B , if A is a subtype of B (upcasting). As illustrated by Rule **TMTMABS**, we omit freshness conditions by adopting the Barendregt convention [1]. Finally, Rule **TMHAND** gives typing for handlers. It requires that the right-hand sides of the return clause and all operation clauses have the same computation type $(B ! \Delta)$, and that all operations mentioned are part of the top-level signature Σ .⁴ The result type takes the form $A ! \Delta \cup \mathcal{O} \Rightarrow B ! \Delta$, capturing the intended handler semantics: given a computation of type $A ! \Delta \cup \mathcal{O}$, the handler (a) produces a result of type B , (b) handles operations \mathcal{O} , and (c) propagates unhandled operations Δ to the output.

⁴ We capture all defined operations along with their types in a global signature Σ .

typing environment $\Gamma ::= \epsilon \mid \Gamma, \varsigma \mid \Gamma, \alpha : \tau \mid \Gamma, \delta \mid \Gamma, x : S \mid \Gamma, \omega : \pi$

$\Gamma \vdash_v v : A \rightsquigarrow v'$ **Values**

$$\frac{(x : \forall \zeta. \forall \alpha : \bar{\tau}. \forall \delta. \bar{\pi} \Rightarrow A) \in \Gamma \quad \sigma = [\bar{\tau}'/\zeta, \bar{B}/\alpha, \bar{\Delta}/\delta] \quad \overline{\Gamma \vdash_{\text{co}} \gamma : \sigma(\pi)}}{\Gamma \vdash_v x : \sigma(A) \rightsquigarrow x \quad \bar{\tau}' \quad \bar{B} \quad \bar{\Delta} \quad \bar{\gamma}} \text{TMVAR}$$

$$\frac{\Gamma \vdash_v v : A \rightsquigarrow v' \quad \overline{\Gamma \vdash_{\text{co}} \gamma : A \leq B}}{\Gamma \vdash_v v : B \rightsquigarrow v' \triangleright \gamma} \text{TMCASTV} \quad \frac{}{\Gamma \vdash_v \text{unit} : \text{Unit} \rightsquigarrow \text{unit}} \text{TMUNIT}$$

$$\frac{\Gamma, x : A \vdash_c c : \underline{C} \rightsquigarrow c' \quad \Gamma \vdash_{\text{vty}} A : \tau \rightsquigarrow T}{\Gamma \vdash_v (\text{fun } x \mapsto c) : A \rightarrow \underline{C} \rightsquigarrow \text{fun } (x : T) \mapsto c'} \text{TMABS}$$

$$\frac{\begin{array}{l} \Gamma, x : A \vdash_c c_r : B ! \Delta \rightsquigarrow c'_r \quad \Gamma \vdash_{\text{vty}} A : \tau \rightsquigarrow T \\ \left[(\text{Op} : A_{\text{Op}} \rightarrow B_{\text{Op}}) \in \Sigma \quad \Gamma, x : A_{\text{Op}}, k : B_{\text{Op}} \rightarrow B ! \Delta \vdash_c c_{\text{Op}} : B ! \Delta \rightsquigarrow c'_{\text{Op}} \right]_{\text{Op} \in \mathcal{O}} \\ c_{\text{res}} = \{ \text{return } (x : T) \mapsto c'_r, [\text{Op } x \ k \mapsto c'_{\text{Op}}]_{\text{Op} \in \mathcal{O}} \} \end{array}}{\Gamma \vdash_v \{ \text{return } x \mapsto c_r, [\text{Op } x \ k \mapsto c_{\text{Op}}]_{\text{Op} \in \mathcal{O}} \} : A ! \Delta \cup \mathcal{O} \Rightarrow B ! \Delta \rightsquigarrow c_{\text{res}}} \text{TMHAND}$$

$\Gamma \vdash_c c : \underline{C} \rightsquigarrow c'$ **Computations**

$$\frac{\Gamma \vdash_c c : \underline{C}_1 \rightsquigarrow c' \quad \overline{\Gamma \vdash_{\text{co}} \gamma : \underline{C}_1 \leq \underline{C}_2}}{\Gamma \vdash_c c : \underline{C}_2 \rightsquigarrow c' \triangleright \gamma} \text{TMCASTC} \quad \frac{\Gamma \vdash_v v_1 : A \rightarrow \underline{C} \rightsquigarrow v'_1 \quad \Gamma \vdash_v v_2 : A \rightsquigarrow v'_2}{\Gamma \vdash_c v_1 \ v_2 : \underline{C} \rightsquigarrow v'_1 \ v'_2} \text{TMAPP}$$

$$\frac{S = \forall \zeta. \forall \alpha : \bar{\tau}. \forall \delta. \bar{\pi} \Rightarrow A \quad \Gamma, \bar{\zeta}, \bar{\alpha} : \bar{\tau}, \bar{\delta}, \bar{\omega} : \bar{\pi} \vdash_v v : A \rightsquigarrow v' \quad \Gamma, x : S \vdash_c c : \underline{C} \rightsquigarrow c'}{\Gamma \vdash_c \text{let } x = v \text{ in } c : \underline{C} \rightsquigarrow \text{let } x = \Lambda \bar{\zeta}. \Lambda \bar{\alpha} : \bar{\tau}. \Lambda \bar{\delta}. \Lambda (\bar{\omega} : \bar{\pi}). v' \text{ in } c'} \text{TMLET}$$

$$\frac{\Gamma \vdash_v v : A \rightsquigarrow v'}{\Gamma \vdash_c \text{return } v : A ! \emptyset \rightsquigarrow \text{return } v'} \text{TMRETURN}$$

$$\frac{(\text{Op} : A_{\text{Op}} \rightarrow B_{\text{Op}}) \in \Sigma \quad \Gamma \vdash_v v : A_{\text{Op}} \rightsquigarrow v' \quad \Gamma, y : B_{\text{Op}} \vdash_c c : A ! \Delta \rightsquigarrow c' \quad \Gamma \vdash_{\text{vty}} B_{\text{Op}} : \tau \rightsquigarrow T_{\text{Op}} \quad \text{Op} \in \Delta}{\Gamma \vdash_c \text{Op } v (y.c) : A ! \Delta \rightsquigarrow \text{Op } v' (y : T_{\text{Op}}.c')} \text{TMOp}$$

$$\frac{\Gamma \vdash_c c_1 : A ! \Delta \rightsquigarrow c'_1 \quad \Gamma, x : A \vdash_c c_2 : B ! \Delta \rightsquigarrow c'_2}{\Gamma \vdash_c \text{do } x \leftarrow c_1; c_2 : B ! \Delta \rightsquigarrow \text{do } x \leftarrow c'_1; c'_2} \text{TMDO}$$

$$\frac{\Gamma \vdash_v v : \underline{C} \Rightarrow \underline{D} \rightsquigarrow v' \quad \Gamma \vdash_c c : \underline{C} \rightsquigarrow c'}{\Gamma \vdash_c \text{handle } c \text{ with } v : \underline{D} \rightsquigarrow \text{handle } c' \text{ with } v'} \text{TMHANDLE}$$

Fig. 2. IMPEFF Typing & Elaboration

Computations. Typing for computations takes the form $\Gamma \vdash_c c : \underline{C} \rightsquigarrow c'$, and, given a typing environment Γ , checks a computation c against a type \underline{C} .

Rule **TMCASTC** behaves like Rule **TMCASTV**, but for computation types. Rule **TMLET** handles polymorphic, non-recursive let-bindings. Rule **TMRETURN** handles **return** v computations. Keyword **return** effectively lifts a value v of type A into a computation of type $A ! \emptyset$. Rule **TMOP** checks operation calls. First, we ensure that v has the appropriate type, as specified by the signature of **Op**. Then, the continuation $(y.c)$ is checked. The side condition $\text{Op} \in \Delta$ ensures that the called operation **Op** is captured in the result type. Rule **TMDO** handles sequencing. Given that c_1 has type $A ! \Delta$, the pure part of the result of type A is bound to term variable x , which is brought in scope for checking c_2 . As we mentioned in Sect. 2, all computations in a **do**-construct should have the same effect set, Δ . Rule **TMHANDLE** eliminates handler types, just as Rule **TMTMAPP** eliminates arrow types.

Constraint Entailment. The specification of constraint entailment takes the form $\Gamma \vdash_{\text{co}} \gamma : \rho$ and is presented in Fig. 3. Notice that we use ρ instead of π , which allows us to capture subtyping between two value types, computation types or dirts, within the same relation. Subtyping can be established in several ways:

Rule **COVAR** handles given assumptions. Rules **VCOREFL** and **DCOREFL** express that subtyping is reflexive, for both value types and dirts. Notice that we do not have a rule for the reflexivity of computation types since, as we illustrate below, it can be established using the reflexivity of their subparts. Rules **VCOTRANS**, **CCOTRANS** and **DCOTRANS** express the transitivity of subtyping for value types, computation types and dirts, respectively. Rule **VCOARR** establishes inequality of arrow types. As usual, the arrow type constructor is contravariant in the argument type. Rules **VCOARRL** and **CCOARRR** are the inversions of Rule **VCOARR**, allowing us to establish the relation between the subparts of the arrow types. Rules **VCOHAND**, **CCOHL**, and **CCOHR** work similarly, for handler types. Rule **CCOCOMP** captures the covariance of type constructor (!), establishing subtyping between two computation types if subtyping is established for their respective subparts. Rules **VCOPURE** and **DCOIMPURE** are its inversions. Finally, Rules **DCONIL** and **DCOOP** establish subtyping between dirts. Rule **DCONIL** captures that the empty dirty set \emptyset is a subdirt of any dirt Δ and Rule **DCOOP** expresses that dirt subtyping preserved under extension with the same operation **Op**.

Well-Formedness of Types, Constraints, Dirts, and Skeletons. The relations $\Gamma \vdash_{\text{vty}} A : \tau \rightsquigarrow T$ and $\Gamma \vdash_{\text{cty}} \underline{C} : \tau \rightsquigarrow \underline{C}$ check the well-formedness of value and computation types respectively. Similarly, relations $\Gamma \vdash_{\text{ct}} \rho \rightsquigarrow \rho$ and $\Gamma \vdash_{\Delta} \Delta$ check the well-formedness of constraints and dirts, respectively.

<div style="border: 1px solid black; padding: 2px; display: inline-block; margin-right: 10px;">$\Gamma \vdash_{\text{co}} \gamma : \rho$</div> Constraint Entailment	
$\frac{(\omega : \pi) \in \Gamma}{\Gamma \vdash_{\text{co}} \omega : \pi} \text{CoVAR}$	$\frac{\Gamma \vdash_{\text{vty}} A : \tau \rightsquigarrow T}{\Gamma \vdash_{\text{co}} \langle T \rangle : A \leq A} \text{VCoREFL}$
$\frac{\Gamma \vdash_{\text{b}} \Delta}{\Gamma \vdash_{\text{co}} \langle \Delta \rangle : \Delta \leq \Delta} \text{DCoREFL}$	$\frac{\Gamma \vdash_{\text{co}} \gamma_1 : A_1 \leq A_2 \quad \Gamma \vdash_{\text{co}} \gamma_2 : A_2 \leq A_3}{\Gamma \vdash_{\text{co}} \gamma_1 \gg \gamma_2 : A_1 \leq A_3} \text{VCoTRANS}$
$\frac{\Gamma \vdash_{\text{co}} \gamma_1 : \underline{C}_1 \leq \underline{C}_2 \quad \Gamma \vdash_{\text{co}} \gamma_2 : \underline{C}_2 \leq \underline{C}_3}{\Gamma \vdash_{\text{co}} \gamma_1 \gg \gamma_2 : \underline{C}_1 \leq \underline{C}_3} \text{CCoTRANS}$	$\frac{\Gamma \vdash_{\text{co}} \gamma_1 : \underline{\Delta}_1 \leq \underline{\Delta}_2 \quad \Gamma \vdash_{\text{co}} \gamma_2 : \underline{\Delta}_2 \leq \underline{\Delta}_3}{\Gamma \vdash_{\text{co}} \gamma_1 \gg \gamma_2 : \underline{\Delta}_1 \leq \underline{\Delta}_3} \text{DCoTRANS}$
$\frac{\Gamma \vdash_{\text{co}} \gamma_1 : B \leq A \quad \Gamma \vdash_{\text{co}} \gamma_2 : \underline{C} \leq \underline{D}}{\Gamma \vdash_{\text{co}} \gamma_1 \rightarrow \gamma_2 : A \rightarrow \underline{C} \leq B \rightarrow \underline{D}} \text{VCoARR}$	
$\frac{\Gamma \vdash_{\text{co}} \gamma : A \rightarrow \underline{C} \leq B \rightarrow \underline{D}}{\Gamma \vdash_{\text{co}} \text{left}(\gamma) : B \leq A} \text{VCoARRL}$	$\frac{\Gamma \vdash_{\text{co}} \gamma : A \rightarrow \underline{C} \leq B \rightarrow \underline{D}}{\Gamma \vdash_{\text{co}} \text{right}(\gamma) : \underline{C} \leq \underline{D}} \text{CCoARRR}$
$\frac{\Gamma \vdash_{\text{co}} \gamma_1 : \underline{C}_2 \leq \underline{C}_1 \quad \Gamma \vdash_{\text{co}} \gamma_2 : \underline{D}_1 \leq \underline{D}_2}{\Gamma \vdash_{\text{co}} \gamma_1 \Rightarrow \gamma_2 : \underline{C}_1 \Rightarrow \underline{D}_1 \leq \underline{C}_2 \Rightarrow \underline{D}_2} \text{VCoHAND}$	
$\frac{\Gamma \vdash_{\text{co}} \gamma : \underline{C}_1 \Rightarrow \underline{D}_1 \leq \underline{C}_2 \Rightarrow \underline{D}_2}{\Gamma \vdash_{\text{co}} \text{left}(\gamma) : \underline{C}_2 \leq \underline{C}_1} \text{CCoHL}$	$\frac{\Gamma \vdash_{\text{co}} \gamma : \underline{C}_1 \Rightarrow \underline{D}_1 \leq \underline{C}_2 \Rightarrow \underline{D}_2}{\Gamma \vdash_{\text{co}} \text{right}(\gamma) : \underline{D}_1 \leq \underline{D}_2} \text{CCoHR}$
$\frac{\Gamma \vdash_{\text{co}} \gamma_1 : A_1 \leq A_2 \quad \Gamma \vdash_{\text{co}} \gamma_2 : \Delta_1 \leq \Delta_2}{\Gamma \vdash_{\text{co}} \gamma_1 ! \gamma_2 : A_1 ! \Delta_1 \leq A_2 ! \Delta_2} \text{CCoCOMP}$	
$\frac{\Gamma \vdash_{\text{co}} \gamma : A_1 ! \Delta_1 \leq A_2 ! \Delta_2}{\Gamma \vdash_{\text{co}} \text{pure}(\gamma) : A_1 \leq A_2} \text{VCoPURE}$	$\frac{\Gamma \vdash_{\text{co}} \gamma : A_1 ! \Delta_1 \leq A_2 ! \Delta_2}{\Gamma \vdash_{\text{co}} \text{impure}(\gamma) : \Delta_1 \leq \Delta_2} \text{DCoIMPURE}$
$\frac{}{\Gamma \vdash_{\text{co}} \emptyset_{\Delta} : \emptyset \leq \Delta} \text{DCoNIL}$	$\frac{\Gamma \vdash_{\text{co}} \gamma : \Delta_1 \leq \Delta_2 \quad (\text{Op} : A_{\text{Op}} \rightarrow B_{\text{Op}}) \in \Sigma}{\Gamma \vdash_{\text{co}} \{\text{Op}\} \cup \gamma : \{\text{Op}\} \cup \Delta_1 \leq \{\text{Op}\} \cup \Delta_2} \text{DCoOP}$

Fig. 3. IMPEFF Constraint Entailment

4 The ExEff Language

4.1 Syntax

Figure 4 presents EXEFF's syntax. EXEFF is an intensional type theory akin to System F [7], where every term encodes its own typing derivation. In essence, all abstractions and applications that are implicit in IMPEFF, are made explicit in EXEFF via new syntactic forms. Additionally, EXEFF is impredicative, which is reflected in the lack of discrimination between value types, qualified types and

Terms

$$\begin{aligned}
\text{value } v &::= x \mid \mathbf{unit} \mid \mathbf{fun} (x : T) \mapsto c \mid h \\
&\quad \mid \Lambda \varsigma.v \mid v \tau \mid \Lambda \alpha : \tau.v \mid v T \mid \Lambda \delta.v \mid v \Delta \mid \Lambda (\omega : \pi).v \mid v \gamma \mid v \triangleright \gamma \\
\text{handler } h &::= \{\mathbf{return} (x : T) \mapsto c_r, \mathbf{Op}_1 x k \mapsto c_{\mathbf{Op}_1}, \dots, \mathbf{Op}_n x k \mapsto c_{\mathbf{Op}_n}\} \\
\text{computation } c &::= \mathbf{return} v \mid \mathbf{Op} v (y : T.c) \mid \mathbf{do} x \leftarrow c_1; c_2 \\
&\quad \mid \mathbf{handle} c \text{ with } v \mid v_1 v_2 \mid \mathbf{let} x = v \text{ in } c \mid c \triangleright \gamma
\end{aligned}$$
Types

$$\begin{aligned}
\text{skeleton } \tau &::= \varsigma \mid \mathbf{Unit} \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \Rightarrow \tau_2 \mid \forall \varsigma.\tau \\
\text{value type } T &::= \alpha \mid \mathbf{Unit} \mid T \rightarrow \underline{C} \mid \underline{C}_1 \Rightarrow \underline{C}_2 \mid \forall \varsigma.T \mid \forall \alpha : \tau.T \mid \forall \delta.T \mid \pi \Rightarrow T \\
\text{simple coercion type } \pi &::= T_1 \leq T_2 \mid \Delta_1 \leq \Delta_2 \\
\text{coercion type } \rho &::= \pi \mid \underline{C}_1 \leq \underline{C}_2 \\
\text{computation type } \underline{C} &::= T \mid \Delta \\
\text{dirt } \Delta &::= \delta \mid \emptyset \mid \{\mathbf{Op}\} \cup \Delta
\end{aligned}$$
Coercions

$$\begin{aligned}
\gamma &::= \omega \mid \gamma_1 \gg \gamma_2 \mid \langle T \rangle \mid \gamma_1 \rightarrow \gamma_2 \mid \gamma_1 \Rightarrow \gamma_2 \mid \mathbf{left}(\gamma) \mid \mathbf{right}(\gamma) \mid \langle \Delta \rangle \mid \emptyset_\Delta \mid \{\mathbf{Op}\} \cup \gamma \\
&\quad \mid \forall \varsigma.\gamma \mid \gamma[\tau] \mid \forall \alpha.\gamma \mid \gamma[T] \mid \forall \delta.\gamma \mid \gamma[\Delta] \mid \pi \Rightarrow \gamma \mid \gamma_1 @ \gamma_2 \mid \gamma_1 ! \gamma_2 \mid \mathbf{pure}(\gamma) \mid \mathbf{impure}(\gamma)
\end{aligned}$$
Fig. 4. EXEFF Syntax

type schemes; all non-computation types are denoted by T . While the impredicativity is not strictly required for the purpose at hand, it makes for a cleaner system.

Coercions. Of particular interest is the use of explicit *subtyping coercions*, denoted by γ . EXEFF uses these to replace the implicit casts of IMPEFF (Rules TMCASVT and TMCASCTC in Fig. 2) with explicit casts ($v \triangleright \gamma$) and ($c \triangleright \gamma$).

Essentially, coercions γ are explicit witnesses of subtyping derivations: each coercion form corresponds to a subtyping rule. Subtyping forms a partial order, which is reflected in coercion forms $\gamma_1 \gg \gamma_2$, $\langle T \rangle$, and $\langle \Delta \rangle$. Coercion form $\gamma_1 \gg \gamma_2$ captures transitivity, while forms $\langle T \rangle$ and $\langle \Delta \rangle$ capture reflexivity for value types and dirts (reflexivity for computation types can be derived from these).

Subtyping for skeleton abstraction, type abstraction, dirt abstraction, and qualification is witnessed by forms $\forall \varsigma.\gamma$, $\forall \alpha.\gamma$, $\forall \delta.\gamma$, and $\pi \Rightarrow \gamma$, respectively. Similarly, forms $\gamma[\tau]$, $\gamma[T]$, $\gamma[\Delta]$, and $\gamma_1 @ \gamma_2$ witness subtyping of skeleton instantiation, type instantiation, dirt instantiation, and coercion application, respectively.

Syntactic forms $\gamma_1 \rightarrow \gamma_2$ and $\gamma_1 \Rightarrow \gamma_2$ capture injection for the arrow and the handler type constructor, respectively. Similarly, inversion forms $\mathbf{left}(\gamma)$ and $\mathbf{right}(\gamma)$ capture projection, following from the injectivity of both type constructors.

Coercion form $\gamma_1 ! \gamma_2$ witnesses subtyping for computation types, using proofs for their components. Inversely, syntactic forms $pure(\gamma)$ and $impure(\gamma)$ witness subtyping between the value- and dirt-components of a computation coercion.

Finally, coercion forms \emptyset_Δ and $\{\mathbf{Op}\} \cup \gamma$ are concerned with dirt subtyping. Form \emptyset_Δ witnesses that the empty dirt \emptyset is a subdirt of any dirt Δ . Lastly, coercion form $\{\mathbf{Op}\} \cup \gamma$ witnesses that subtyping between dirts is preserved under extension with a new operation. Note that we do not have an inversion form to extract a witness for $\Delta_1 \leq \Delta_2$ from a coercion for $\{\mathbf{Op}\} \cup \Delta_1 \leq \{\mathbf{Op}\} \cup \Delta_2$. The reason is that dirt sets are sets and not inductive structures. For instance, for $\Delta_1 = \{\mathbf{Op}\}$ and $\Delta_2 = \emptyset$ the latter subtyping holds, but the former does not.

4.2 Typing

Value and Computation Typing. Typing for EXEFF values and computations is presented in Figs. 5 and 6 and is given by two mutually recursive relations of the form $\Gamma \vdash_v v : T$ (values) and $\Gamma \vdash_c c : \underline{C}$ (computations). EXEFF typing environments Γ contain bindings for variables of all sorts:

$$\Gamma ::= \epsilon \mid \Gamma, \varsigma \mid \Gamma, \alpha : \tau \mid \Gamma, \delta \mid \Gamma, x : T \mid \Gamma, \omega : \pi$$

Typing is entirely syntax-directed. Apart from the typing rules for skeleton, type, dirt, and coercion abstraction (and, subsequently, skeleton, type, dirt, and coercion application), the main difference between typing for IMPEFF and EXEFF lies in the explicit cast forms, $(v \triangleright \gamma)$ and $(c \triangleright \gamma)$. Given that a value v has type T_1 and that γ is a proof that T_1 is a subtype of T_2 , we can upcast v with an explicit cast operation $(v \triangleright \gamma)$. Upcasting for computations works analogously.

$\frac{(x : T) \in \Gamma}{\Gamma \vdash_v x : T}$	$\frac{}{\Gamma \vdash_v \mathbf{unit} : \mathbf{Unit}}$	$\frac{\Gamma, x : T \vdash_c c : \underline{C} \quad \Gamma \vdash_T T : \tau}{\Gamma \vdash_v (\mathbf{fun } x : T \mapsto c) : T \rightarrow \underline{C}}$
$\frac{\Gamma \vdash_v v : T_1 \quad \Gamma \vdash_{\infty} \gamma : T_1 \leq T_2}{\Gamma \vdash_v v \triangleright \gamma : T_2}$	$\frac{\Gamma, \varsigma \vdash_v v : T}{\Gamma \vdash_v \Lambda \varsigma . v : \forall \varsigma . T}$	$\frac{\Gamma, \alpha : \tau \vdash_v v : T}{\Gamma \vdash_v \Lambda \alpha : \tau . v : \forall \alpha : \tau . T}$
$\frac{\Gamma, \delta \vdash_v v : T}{\Gamma \vdash_v \Lambda \delta . v : \forall \delta . T}$	$\frac{\Gamma, \omega : \pi \vdash_v v : T \quad \Gamma \vdash_\rho \pi}{\Gamma \vdash_v \Lambda (\omega : \pi) . v : \pi \Rightarrow T}$	$\frac{\Gamma \vdash_v v : \pi \Rightarrow T \quad \Gamma \vdash_{\infty} \gamma : \pi}{\Gamma \vdash_v v \gamma : T}$
$\frac{\Gamma, x : T_x \vdash_c c_r : T ! \Delta \quad [(\mathbf{Op} : T_1 \rightarrow T_2) \in \Sigma \quad \Gamma, x : T_1, k : T_2 \rightarrow T ! \Delta \vdash_c c_{\mathbf{Op}} : T ! \Delta]_{\mathbf{Op} \in \mathcal{O}}}{\Gamma \vdash_v \{\mathbf{return } (x : T_x) \mapsto c_r, [\mathbf{Op } x k \mapsto c_{\mathbf{Op}}]_{\mathbf{Op} \in \mathcal{O}}\} : T_x ! \Delta \cup \mathcal{O} \Rightarrow T ! \Delta}$		
$\frac{\Gamma \vdash_v v : \forall \varsigma . T \quad \Gamma \vdash_\tau \tau}{\Gamma \vdash_v v \tau : T[\tau/\varsigma]}$	$\frac{\Gamma \vdash_v v : \forall \alpha : \tau . T_1 \quad \Gamma \vdash_T T_2 : \tau}{\Gamma \vdash_v v T_2 : T_1[T_2/\alpha]}$	$\frac{\Gamma \vdash_v v : \forall \delta . T \quad \Gamma \vdash_\delta \Delta}{\Gamma \vdash_v v \Delta : T[\Delta/\delta]}$

Fig. 5. EXEFF Value Typing

Well-Formedness of Types, Constraints, Dirts and Skeletons. The definitions of the judgements that check the well-formedness of EXEFF value types ($\Gamma \vdash_T T : \tau$), computation types ($\Gamma \vdash_{\underline{C}} \underline{C} : \tau$), dirt ($\Gamma \vdash_{\Delta} \Delta$), and skeletons ($\Gamma \vdash_{\tau} \tau$) are equally straightforward as those for IMPEFF.

Coercion Typing. Coercion typing formalizes the intuitive interpretation of coercions we gave in Sect. 4.1 and takes the form $\Gamma \vdash_{\circ} \gamma : \rho$. It is essentially an extension of the constraint entailment relation of Fig. 3.

4.3 Operational Semantics

Figure 7 presents selected rules of EXEFF's small-step, call-by-value operational semantics. For lack of space, we omit β -rules and other common rules and focus only on cases of interest.

Firstly, one of the non-conventional features of our system lies in the stratification of results in plain results and cast results:

$$\begin{array}{c}
 \frac{\Gamma \vdash_v v_1 : T \rightarrow \underline{C} \quad \Gamma \vdash_v v_2 : T}{\Gamma \vdash_c v_1 v_2 : \underline{C}} \qquad \frac{\Gamma \vdash_v v : T \quad \Gamma, x : T \vdash_c c : \underline{C}}{\Gamma \vdash_c \text{let } x = v \text{ in } c : \underline{C}} \\
 \\
 \frac{\Gamma \vdash_c v : T}{\Gamma \vdash_c \text{return } v : T ! \emptyset} \qquad \frac{\Gamma \vdash_c c_1 : T_1 ! \Delta \quad \Gamma, x : T_1 \vdash_c c_2 : T_2 ! \Delta}{\Gamma \vdash_c \text{do } x \leftarrow c_1; c_2 : T_2 ! \Delta} \\
 \\
 \frac{(\text{Op} : T_1 \rightarrow T_2) \in \Sigma \quad \Gamma \vdash_v v : T_1 \quad \Gamma, y : T_2 \vdash_c c : T ! \Delta \quad \text{Op} \in \Delta}{\Gamma \vdash_c \text{Op } v (y : T_2.c) : T ! \Delta} \\
 \\
 \frac{\Gamma \vdash_v v : \underline{C}_1 \Rightarrow \underline{C}_2 \quad \Gamma \vdash_c c : \underline{C}_1}{\Gamma \vdash_c \text{handle } c \text{ with } v : \underline{C}_2} \qquad \frac{\Gamma \vdash_c c : \underline{C}_1 \quad \Gamma \vdash_{\circ} \gamma : \underline{C}_1 \leq \underline{C}_2}{\Gamma \vdash_c c \triangleright \gamma : \underline{C}_2}
 \end{array}$$

Fig. 6. EXEFF Computation Typing

$$\begin{array}{l}
 \text{terminal value } v^T ::= \text{unit} \mid h \mid \text{fun } x : T \mapsto c \mid \Lambda \alpha : \tau.v \mid \Lambda \delta.v \mid \lambda \omega : \pi.v \\
 \text{value result } v^R ::= v^T \mid v^T \triangleright \gamma \\
 \text{computation result } c^R ::= \text{return } v^T \mid (\text{return } v^T) \triangleright \gamma \mid \text{Op } v^R (y : T.c)
 \end{array}$$

Terminal values v^T represent conventional values, and value results v^R can either be plain terminal values v^T or terminal values with a cast: $v^T \triangleright \gamma$. The same applies to computation results c^R .⁵

Although unusual, this stratification can also be found in Crary's coercion calculus for inclusive subtyping [4], and, more recently, in System F_C [25]. Stratification is crucial for ensuring type preservation. Consider for example the expression

⁵ Observe that operation values do not feature an outermost cast operation, as the coercion can always be pushed into its continuation.

(`return` 5 \triangleright $\langle \text{int} \rangle ! \emptyset_{\{\text{Op}\}}$), of type $\text{int} ! \{\text{Op}\}$. We can not reduce the expression further without losing effect information; removing the cast would result in computation (`return` 5), of type $\text{int} ! \emptyset$. Even if we consider type preservation only up to subtyping, the redex may still occur as a subterm in a context that expects solely the larger type.

Secondly, we need to make sure that casts do not stand in the way of evaluation. This is captured in the so-called “push” rules, all of which appear in Fig. 7.

In relation $v \rightsquigarrow_v v'$, the first rule groups nested casts into a single cast, by means of transitivity. The next three rules capture the essence of push rules: whenever a redex is “blocked” due to a cast, we take the coercion apart and redistribute it (in a type-preserving manner) over the subterms, so that evaluation can progress.

The situation in relation $c \rightsquigarrow_c c'$ is quite similar. The first rule uses transitivity to group nested casts into a single cast. The second rule is a push rule for β -reduction. The third rule pushes a cast out of a `return`-computation. The fourth rule pushes a coercion inside an operation-computation, illustrating why the syntax for c^R does not require casts on operation-computations. The fifth rule is a push rule for sequencing computations and performs two tasks at once. Since we know that the computation bound to x calls no operations, we (a) safely “drop” the impure part of γ , and (b) substitute x with v^T , cast with the pure part of γ (so that types are preserved). The sixth rule handles operation calls in sequencing computations. If an operation is called in a sequencing computation, evaluation is suspended and the rest of the computation is captured in the continuation.

The last four rules are concerned with effect handling. The first of them pushes a coercion on the handler “outwards”, such that the handler can be exposed and evaluation is not stuck (similarly to the push rule for term application). The second rule behaves similarly to the push/beta rule for sequencing computations. Finally, the last two rules are concerned with handling of operations. The first of the two captures cases where the called operation is handled by the handler, in which case the respective clause of the handler is called. As illustrated by the rule, like Pretnar [20], EXEFF features *deep handlers*: the continuation is also wrapped within a `with-handle` construct. The last rule captures cases where the operation is not covered by the handler and thus remains unhandled.

We have shown that EXEFF is type safe:

Theorem 1 (Type Safety)

- If $\Gamma \vdash_v v : T$ then either v is a result value or $v \rightsquigarrow_v v'$ and $\Gamma \vdash_v v' : T$.
- If $\Gamma \vdash_c c : \underline{C}$ then either c is a result computation or $c \rightsquigarrow_c c'$ and $\Gamma \vdash_c c' : \underline{C}$.

$v \rightsquigarrow_v v'$ **Values**

$$(v^T \triangleright \gamma_1) \triangleright \gamma_2 \rightsquigarrow_v v^T \triangleright (\gamma_1 \gg \gamma_2) \quad (v^T \triangleright \gamma) T \rightsquigarrow_v (v^T T) \triangleright \gamma[T]$$

$$(v^T \triangleright \gamma) \Delta \rightsquigarrow_v (v^T \Delta) \triangleright \gamma[\Delta] \quad (v^T \triangleright \gamma_1) \gamma_2 \rightsquigarrow_v (v^T \gamma_2) \triangleright \gamma_1 @ \gamma_2$$

 $c \rightsquigarrow_c c'$ **Computations**

$$(c^R \triangleright \gamma_1) \triangleright \gamma_2 \rightsquigarrow_c c^R \triangleright (\gamma_1 \gg \gamma_2) \quad (v_1^T \triangleright \gamma) v_2 \rightsquigarrow_c (v_1^T (v_2 \triangleright \text{left}(\gamma))) \triangleright \text{right}(\gamma)$$

$$\text{return } (v^T \triangleright \gamma) \rightsquigarrow_c (\text{return } v^T) \triangleright (\gamma ! \emptyset_\emptyset)$$

$$(\text{Op } v^R (y : T.c)) \triangleright \gamma \rightsquigarrow_c \text{Op } v^R (y : T.(c \triangleright \gamma))$$

$$\text{do } x \leftarrow ((\text{return } v^T) \triangleright \gamma); c_2 \rightsquigarrow_c c_2[(v^T \triangleright \text{pure}(\gamma))/x]$$

$$\text{do } x \leftarrow \text{Op } v^R (y : T.c_1); c_2 \rightsquigarrow_c \text{Op } v^R (y : T.\text{do } x \leftarrow c_1; c_2)$$

$$\text{handle } c \text{ with } (v^T \triangleright \gamma) \rightsquigarrow_c (\text{handle } (c \triangleright \text{left}(\gamma)) \text{ with } v^T) \triangleright \text{right}(\gamma)$$

$$\text{handle } ((\text{return } v^T) \triangleright \gamma) \text{ with } h \rightsquigarrow_c c_r[v^T \triangleright \text{pure}(\gamma)/x]$$

$$\text{handle } (\text{Op } v^R (y : T.c)) \text{ with } h \rightsquigarrow_c c_{\text{op}}[v^R/x, (\text{fun } (y : T) \mapsto \text{handle } c \text{ with } h)/k]$$

$$\text{handle } (\text{Op } v^R (y : T.c)) \text{ with } h \rightsquigarrow_c \text{Op } v^R (y : T.\text{handle } c \text{ with } h)$$

Fig. 7. EXEFF Operational Semantics (Selected Rules)

5 Type Inference and Elaboration

This section presents the typing-directed elaboration of IMPEFF into EXEFF. This elaboration makes all the implicit type and effect information explicit, and introduces explicit term-level coercions to witness the use of subtyping.

After covering the declarative specification of this elaboration, we present a constraint-based algorithm to infer IMPEFF types and at the same time elaborate into EXEFF. This algorithm alternates between two phases: (1) the syntax-directed generation of constraints from the IMPEFF term, and (2) solving these constraints.

5.1 Elaboration of ImpEff into ExEff

The grayed parts of Fig. 2 augment the typing rules for IMPEFF value and computation terms with typing-directed elaboration to corresponding EXEFF terms. The elaboration is mostly straightforward, mapping every IMPEFF construct onto its corresponding EXEFF construct while adding explicit type annotations to binders in Rules TMTMABS, TMHANDLER and TMOP. Implicit appeals to

subtyping are turned into explicit casts with coercions in Rules `TMCASTV` and `TMCASTC`. Rule `TMLET` introduces explicit binders for skeleton, type, and dirt variables, as well as for constraints. These last also introduce coercion variables ω that can be used in casts. The binders are eliminated in rule `TMVAR` by means of explicit application with skeletons, types, dirts and coercions. The coercions are produced by the auxiliary judgement $\Gamma \vdash_{\text{co}} \gamma : \pi$, defined in Fig. 3, which provides a coercion witness for every subtyping proof.

As a sanity check, we have shown that elaboration preserves types.

Theorem 2 (Type Preservation)

- If $\Gamma \vdash_v v : A \rightsquigarrow v'$ then $\text{elab}_R(\Gamma) \vdash_v v' : \text{elab}_S(A)$.
- If $\Gamma \vdash_c c : \underline{C} \rightsquigarrow c'$ then $\text{elab}_R(\Gamma) \vdash_c c' : \text{elab}_{\underline{C}}(\underline{C})$.

Here $\text{elab}_R(\Gamma)$, $\text{elab}_S(A)$ and $\text{elab}_{\underline{C}}(\underline{C})$ convert `IMPEFF` environments and types into `EXEFF` environments and types.

5.2 Constraint Generation and Elaboration

Constraint generation with elaboration into `EXEFF` is presented in Figs. 8 (values) and 9 (computations). Before going into the details of each, we first introduce the three auxiliary constructs they use.

constraint set $\mathcal{P}, \mathcal{Q} ::= \bullet \mid \tau_1 = \tau_2, \mathcal{P} \mid \alpha : \tau, \mathcal{P} \mid \omega : \pi, \mathcal{P}$
 typing environment $\Gamma ::= \epsilon \mid \Gamma, x : S$
 substitution $\sigma ::= \bullet \mid \sigma \cdot [\tau/s] \mid \sigma \cdot [A/\alpha] \mid \sigma \cdot [\Delta/\delta] \mid \sigma \cdot [\gamma/\omega]$

At the heart of our algorithm are sets \mathcal{P} , containing three different kinds of constraints: (a) skeleton equalities of the form $\tau_1 = \tau_2$, (b) skeleton constraints of the form $\alpha : \tau$, and (c) wanted subtyping constraints of the form $\omega : \pi$. The purpose of the first two becomes clear when we discuss constraint solving, in Sect. 5.3. Next, typing environments Γ only contain term variable bindings, while other variables represent unknowns of their sort and may end up being instantiated after constraint solving. Finally, during type inference we compute substitutions σ , for refining as of yet unknown skeletons, types, dirts, and coercions. The last one is essential, since our algorithm simultaneously performs type inference and elaboration into `EXEFF`.

A substitution σ is a solution of the set \mathcal{P} , written as $\sigma \models \mathcal{P}$, if we get derivable judgements after applying σ to all constraints in \mathcal{P} .

Values. Constraint generation for values takes the form $\mathcal{Q}; \Gamma \vdash_v v : A \mid \mathcal{Q}'; \sigma \rightsquigarrow v'$. It takes as inputs a set of wanted constraints \mathcal{Q} , a typing environment Γ , and a `IMPEFF` value v , and produces a value type A , a new set of wanted constraints \mathcal{Q}' , a substitution σ , and a `EXEFF` value v' .

Unlike standard HM, our inference algorithm does not keep constraint generation and solving separate. Instead, the two are interleaved, as indicated by

$\mathcal{Q}; \Gamma \vdash v : A \mid \mathcal{Q}'; \sigma \rightsquigarrow v'$	Values
$\frac{(x : \forall \bar{\zeta}. \bar{\alpha} : \bar{\tau}. \forall \bar{\delta}. \bar{\pi} \Rightarrow A) \in \Gamma \quad \sigma = [\bar{\zeta}'/\bar{\zeta}, \bar{\alpha}'/\bar{\alpha}, \bar{\delta}'/\bar{\delta}]}{\mathcal{Q}; \Gamma \vdash x : \sigma(A) \mid \bar{\omega} : \sigma(\pi), \alpha' : \sigma(\tau), \mathcal{Q}; \bullet \rightsquigarrow x \zeta' \bar{\alpha}' \bar{\delta}' \bar{\omega}}$	
$\frac{}{\mathcal{Q}; \Gamma \vdash \text{unit} : \text{Unit} \mid \mathcal{Q}; \bullet \rightsquigarrow \text{unit}}$	
$\frac{\alpha : \varsigma, \mathcal{Q}; \Gamma, x : \alpha \vdash c : \underline{C} \mid \mathcal{Q}'; \sigma \rightsquigarrow c'}{\mathcal{Q}; \Gamma \vdash (\text{fun } x \mapsto c) : \sigma(\alpha) \rightarrow \underline{C} \mid \mathcal{Q}'; \sigma \rightsquigarrow \text{fun } x : \sigma(\alpha) \mapsto c'}$	
$\alpha_r : \varsigma_r, \mathcal{Q}; \Gamma, x : \alpha_r \vdash c_r : B_r \mid \Delta_r \mid \mathcal{Q}_0; \sigma_r \rightsquigarrow c'_r \quad \sigma^i = \sigma_i \cdot \sigma_{i-1} \cdot \dots \cdot \sigma_1$	
$\text{Op}_i \in \mathcal{O} :$	$(\text{Op}_i : A_i \rightarrow B_i) \in \Sigma$
$\alpha_i : \varsigma_i, \mathcal{Q}_{i-1}; \sigma^{i-1}(\sigma_r(\Gamma)), x : A_i, k : B_i \rightarrow \alpha_i ! \delta_i \vdash c_{\text{Op}_i} : B_{\text{Op}_i} ! \Delta_{\text{Op}_i} \mid \mathcal{Q}_i; \sigma_i \rightsquigarrow c'_{\text{Op}_i}$	$\mathcal{Q}' = \alpha_{in} : \varsigma_{in}, \alpha_{out} : \varsigma_{out}, \bar{\omega}_1 : \sigma^n(B_r) \leq \alpha_{out}, \bar{\omega}_2 : \sigma^n(\Delta_r) \leq \delta_{out}, \bar{\omega}_{3_i} : \sigma^n(B_{\text{Op}_i}) \leq \alpha_{out}$
$\bar{\omega}_{4_i} : \sigma^n(\Delta_{\text{Op}_i}) \leq \delta_{out}$	$\bar{\omega}_{5_i} : B_i \rightarrow \alpha_{out} ! \delta_{out} \leq B_i \rightarrow \sigma^n(\alpha_i ! \delta_i)^n$
$\bar{\omega}_6 : \alpha_{in} \leq \sigma^n(\sigma_r(\alpha_r)), \bar{\omega}_7 : \delta_{in} \leq \delta_{out} \cup \mathcal{O}, \mathcal{Q}_n$	$c_{res} = \{ \text{return } y : \sigma^n(\sigma_r(\alpha_r)) \mapsto \sigma^n(c'_r)[y \triangleright \omega_6/x] \triangleright \omega_1 ! \omega_2$
$\left. , [\text{Op}_i \ x \ l \mapsto \sigma^n(c'_{\text{Op}_i})[l \triangleright \omega_{5_i}/k] \triangleright \omega_{3_i} ! \omega_{4_i}]_{\text{Op}_i \in \mathcal{O}} \right\} \triangleright (\langle \alpha_{in} \rangle ! \omega_7 \Rightarrow \langle \alpha_{out} \rangle ! \langle \delta_{out} \rangle)$	$\mathcal{Q}; \Gamma \vdash \{ \text{return } x \mapsto c_r, [\text{Op } x \ k \mapsto c_{\text{Op}}]_{\text{Op} \in \mathcal{O}} \} : \alpha_{in} ! \delta_{in} \Rightarrow \alpha_{out} ! \delta_{out} \mid \mathcal{Q}'; (\sigma^n \cdot \sigma_r) \rightsquigarrow c_{res}$

Fig. 8. Constraint Generation with Elaboration (Values)

the additional arguments of our relation: (a) constraints \mathcal{Q} are passed around in a stateful manner (i.e., they are input and output), and (b) substitutions σ generated from constraint solving constitute part of the relation output. We discuss the reason for this interleaved approach in Sect. 5.4; we now focus on the algorithm.

The rules are syntax-directed on the input IMPEFF value. The first rule handles term variables x : as usual for constraint-based type inference the rule instantiates the polymorphic type $(\forall \bar{\zeta}. \bar{\alpha} : \bar{\tau}. \forall \bar{\delta}. \bar{\pi} \Rightarrow A)$ of x with fresh variables; these are placeholders that are determined during constraint solving. Moreover, the rule extends the wanted constraints \mathcal{P} with $\bar{\pi}$, appropriately instantiated. In EXEFF, this corresponds to explicit skeleton, type, dirt, and coercion applications.

More interesting is the third rule, for term abstractions. Like in standard Hindley-Damas-Milner [5], it generates a fresh type variable α for the type of the abstracted term variable x . In addition, it generates a fresh skeleton variable ς , to capture the (yet unknown) shape of α .

As explained in detail in Sect. 5.3, the constraint solver instantiates type variables only through their skeletons annotations. Because we want to allow local constraint solving for the body c of the term abstraction the opportunity to

produce a substitution σ that instantiates α , we have to pass in the annotation constraint $\alpha : \zeta$.⁶ We apply the resulting substitution σ to the result type $\sigma(\alpha) \rightarrow \underline{C}$.⁷

Finally, the fourth rule is concerned with handlers. Since it is the most complex of the rules, we discuss each of its premises separately:

Firstly, we infer a type $B_r ! \Delta_r$ for the right hand side of the **return**-clause. Since α_r is a fresh unification variable, just like for term abstraction we require $\alpha_r : \zeta_r$, for a fresh skeleton variable ζ_r .

Secondly, we check every operation clause in \mathcal{O} in order. For each clause, we generate fresh skeleton, type, and dirt variables (ζ_i , α_i , and δ_i), to account for the (yet unknown) result type $\alpha_i ! \delta_i$ of the continuation k , while inferring type $B_{\text{op}_i} ! \Delta_{\text{op}_i}$ for the right-hand-side c_{op_i} .

More interesting is the (final) set of wanted constraints \mathcal{Q}' . First, we assign to the handler the overall type

$$\alpha_{in} ! \delta_{in} \Rightarrow \alpha_{out} ! \delta_{out}$$

where $\zeta_{in}, \alpha_{in}, \delta_{in}, \zeta_{out}, \alpha_{out}, \delta_{out}$ are fresh variables of the respective sorts. In turn, we require that (a) the type of the return clause is a subtype of $\alpha_{out} ! \delta_{out}$ (given by the combination of ω_1 and ω_2), (b) the right-hand-side type of each operation clause is a subtype of the overall result type: $\sigma^n(B_{\text{op}_i} ! \Delta_{\text{op}_i}) \leq \alpha_{out} ! \delta_{out}$ (witnessed by $\omega_3 ! \omega_4$), (c) the actual types of the continuations $B_i \rightarrow \alpha_{out} ! \delta_{out}$ in the operation clauses should be subtypes of their assumed types $B_i \rightarrow \sigma^n(\alpha_i ! \delta_i)$ (witnessed by ω_5). (d) the overall argument type α_{in} is a subtype of the assumed type of x : $\sigma^n(\sigma_r(\alpha_r))$ (witnessed by ω_6), and (e) the input dirt set δ_{in} is a subtype of the resulting dirt set δ_{out} , extended with the handled operations \mathcal{O} (witnessed by ω_7).

All the aforementioned implicit subtyping relations become explicit in the elaborated term c_{res} , via explicit casts.

Computations. The judgement $\mathcal{Q}; \Gamma \vdash_c c : \underline{C} \mid \mathcal{Q}'; \sigma \rightsquigarrow c'$ generates constraints for computations.

The first rule handles term applications of the form $v_1 v_2$. After inferring a type for each subterm (A_1 for v_1 and A_2 for v_2), we generate the wanted constraint $\sigma_2(A_1) \leq A_2 \rightarrow \alpha ! \delta$, with fresh type and dirt variables α and δ , respectively. Associated coercion variable ω is then used in the elaborated term to explicitly (up)cast v'_1 to the expected type $A_2 \rightarrow \alpha ! \delta$.

The third rule handles polymorphic let-bindings. First, we infer a type A for v , as well as wanted constraints \mathcal{Q}_v . Then, we simplify wanted constraints \mathcal{Q}_v by means of function `solve` (which we explain in detail in Sect. 5.3 below), obtaining a substitution σ'_1 and a set of *residual constraints* \mathcal{Q}'_v .

⁶ This hints at why we need to pass constraints in a stateful manner.

⁷ Though σ refers to IMPEFF types, we abuse notation to save clutter and apply it directly to EXEFF entities too.

$\mathcal{Q}; \Gamma \vdash_c c : \underline{C} \mid \mathcal{Q}'; \sigma \rightsquigarrow c'$	Computations
$\mathcal{Q}; \Gamma \vdash_v v_1 : A_1 \mid \mathcal{Q}_1; \sigma_1 \rightsquigarrow v'_1 \quad \mathcal{Q}_1; \sigma_1(\Gamma) \vdash_v v_2 : A_2 \mid \mathcal{Q}_2; \sigma_2 \rightsquigarrow v'_2$	$\mathcal{Q}; \Gamma \vdash_c v_1 v_2 : \alpha ! \delta \mid \alpha : \varsigma, \omega : \sigma_2(A_1) \leq A_2 \rightarrow \alpha ! \delta, \mathcal{Q}_2; (\sigma_2 \cdot \sigma_1) \rightsquigarrow (\sigma_2(v'_1) \triangleright \omega) v'_2$
$\mathcal{Q}; \Gamma \vdash_v v : A \mid \mathcal{Q}'; \sigma \rightsquigarrow v'$	$\mathcal{Q}; \Gamma \vdash_c \text{return } v : A ! \emptyset \mid \mathcal{Q}'; \sigma \rightsquigarrow \text{return } v'$
$\text{solve}(\bullet; \bullet; \mathcal{Q}_v) = (\sigma'_1, \mathcal{Q}'_v) \quad \text{split}(\sigma'_1(\sigma_1(\Gamma)), \mathcal{Q}'_v, \sigma'_1(A)) = \langle \bar{\varsigma}, \bar{\alpha} : \bar{\tau}, \bar{\delta}, \bar{\omega} : \bar{\pi}, \mathcal{Q}_1 \rangle$	$\mathcal{Q}_1; \sigma'_1(\sigma_1(\Gamma)), x : \forall \bar{\varsigma}. \forall \bar{\alpha} : \bar{\tau}. \forall \bar{\delta}. \bar{\pi} \Rightarrow \sigma'_1(A) \vdash_c c : \underline{C} \mid \mathcal{Q}_2; \sigma_2 \rightsquigarrow c'$
$c_{res} = \text{let } x = \sigma_2(\Lambda \bar{\varsigma}. \Lambda \bar{\alpha} : \bar{\tau}. \Lambda \bar{\delta}. \Lambda(\omega : \text{elab}_\rho(\bar{\pi})). v') \text{ in } c'$	$\mathcal{Q}; \Gamma \vdash_c \text{let } x = v \text{ in } c : \underline{C} \mid \mathcal{Q}_2; (\sigma_2 \cdot \sigma'_1 \cdot \sigma_1) \rightsquigarrow c_{res}$
$\mathcal{Q}; \Gamma \vdash_v v : A_1 \mid \mathcal{Q}_1; \sigma_1 \rightsquigarrow v' \quad \mathcal{Q}_1; \sigma_1(\Gamma), y : B_{\text{Op}} \vdash_c c : A_2 ! \Delta_2 \mid \mathcal{Q}_2; \sigma_2 \rightsquigarrow c'$	$(\text{Op} : A_{\text{Op}} \rightarrow B_{\text{Op}}) \in \Sigma \quad c_{res} = \text{Op } (\sigma_2(v') \triangleright \omega) (y : \text{elab}_S(B_{\text{Op}}). c')$
$\mathcal{Q}; \Gamma \vdash_c \text{Op } v (y : B_{\text{Op}}. c) : A_2 ! \{\text{Op}\} \cup \Delta_2 \mid \omega : \sigma_2(A_1) \leq A_{\text{Op}}, \mathcal{Q}_2; (\sigma_2 \cdot \sigma_1) \rightsquigarrow c_{res}$	
$\mathcal{Q}; \Gamma \vdash_c c_1 : A_1 ! \Delta_1 \mid \mathcal{Q}_1; \sigma_1 \rightsquigarrow c'_1 \quad \mathcal{Q}_1; \sigma_1(\Gamma), x : A_1 \vdash_c c_2 : A_2 ! \Delta_2 \mid \mathcal{Q}_2; \sigma_2 \rightsquigarrow c'_2$	$c_{res} = \text{do } x \leftarrow (\sigma_2(c'_1) \triangleright (\sigma_2(A_1))) ! \omega_1; (c'_2 \triangleright (A_2)) ! \omega_2$
$\mathcal{Q}; \Gamma \vdash_c \text{do } x \leftarrow c_1; c_2 : A_2 ! \delta \mid \omega_1 : \sigma_2(\Delta_1) \leq \delta, \omega_2 : \Delta_2 \leq \delta, \mathcal{Q}_2; (\sigma_2 \cdot \sigma_1) \rightsquigarrow c_{res}$	
$\mathcal{Q}; \Gamma \vdash_v v : A_1 \mid \mathcal{Q}_1; \sigma_1 \rightsquigarrow v' \quad \mathcal{Q}_1; \sigma_1(\Gamma) \vdash_c c : A_2 ! \Delta_2 \mid \mathcal{Q}_2; \sigma_2 \rightsquigarrow c'$	$\mathcal{Q}' = \alpha_1 : \varsigma_1, \alpha_2 : \varsigma_2, \omega_1 : \sigma_2(A_1) \leq (\alpha_1 ! \delta_1 \Rightarrow \alpha_2 ! \delta_2), \omega_2 : A_2 \leq \alpha_1, \omega_3 : \Delta_2 \leq \delta_1, \mathcal{Q}_2$
$c_{res} = \text{handle } (c' \triangleright (\omega_2 ! \omega_3)) \text{ with } (\sigma_2(v') \triangleright \omega_1)$	$\mathcal{Q}; \Gamma \vdash_c \text{handle } c \text{ with } v : \alpha_2 ! \Delta_2 \mid \mathcal{Q}'; (\sigma_2 \cdot \sigma_1) \rightsquigarrow c_{res}$

Fig. 9. Constraint Generation with Elaboration (Computations)

Generalization of x 's type is performed by auxiliary function *split*, given by the following clause:

$$\begin{array}{l}
 \bar{\varsigma} = \{\varsigma \mid (\alpha : \varsigma) \in \mathcal{Q}, \nexists \alpha'. \alpha' \notin \bar{\alpha} \wedge (\alpha' : \varsigma) \in \mathcal{Q}\} \\
 \bar{\alpha} = fv_\alpha(\mathcal{Q}) \cup fv_\alpha(A) \setminus fv_\alpha(\Gamma) \quad \mathcal{Q}_1 = \{(\omega : \pi) \mid (\omega : \pi) \in \mathcal{Q}, fv(\pi) \not\subseteq fv(\Gamma)\} \\
 \bar{\delta} = fv_\delta(\mathcal{Q}) \cup fv_\delta(A) \setminus fv_\delta(\Gamma) \quad \mathcal{Q}_2 = \mathcal{Q} - \mathcal{Q}_1 \\
 \hline
 \text{split}(\Gamma, \mathcal{Q}, A) = \langle \bar{\varsigma}, \bar{\alpha} : \bar{\tau}, \bar{\delta}, \mathcal{Q}_1, \mathcal{Q}_2 \rangle
 \end{array}$$

In essence, *split* generates the type (scheme) of x in parts. Additionally, it computes the subset \mathcal{Q}_2 of the input constraints \mathcal{Q} that do not depend on locally-bound variables. Such constraints can be floated “upwards”, and are passed as input when inferring a type for c . The remainder of the rule is self-explanatory.

The fourth rule handles operation calls. Observe that in the elaborated term, we upcast the inferred type to match the expected type in the signature.

The fifth rule handles sequences. The requirement that all computations in a `do`-construct have the same dirt set is expressed in the wanted constraints $\sigma_2(\Delta_1) \leq \delta$ and $\Delta_2 \leq \delta$ (where δ is a fresh dirt variable; the resulting dirt set), witnessed by coercion variables ω_1 and ω_2 . Both coercion variables are used in the elaborated term to upcast c_1 and c_2 , such that both draw effects from the same dirt set δ .

Finally, the sixth rule is concerned with effect handling. After inferring type A_1 for the handler v , we require that it takes the form of a handler type, witnessed by coercion variable $\omega_1 : \sigma_2(A_1) \leq (\alpha_1! \delta_1 \Rightarrow \alpha_2! \delta_2)$, for fresh $\alpha_1, \alpha_2, \delta_1, \delta_2$. To ensure that the type $A_2! \Delta_2$ of c matches the expected type, we require that $A_2! \Delta_2 \leq \alpha_1! \delta_1$. Our syntax does not include coercion variables for computation subtyping; we achieve the same effect by combining $\omega_2 : A_2 \leq \alpha_1$ and $\omega_3 : \Delta_2 \leq \delta_1$.

Theorem 3 (Soundness of Inference). *If $\bullet; \Gamma \vdash_v v : A \mid \mathcal{Q}; \sigma \rightsquigarrow v'$ then for any $\sigma' \models \mathcal{Q}$, we have $(\sigma' \cdot \sigma)(\Gamma) \vdash_v v : \sigma'(A) \rightsquigarrow \sigma'(v')$, and analogously for computations.*

Theorem 4 (Completeness of Inference). *If $\Gamma \vdash_v v : A \rightsquigarrow v'$ then we have $\bullet; \Gamma \vdash_v v : A' \mid \mathcal{Q}; \sigma \rightsquigarrow v''$ and there exists $\sigma' \models \mathcal{Q}$ and γ , such that $\sigma'(v'') = v'$ and $\sigma(\Gamma) \vdash_{\infty} \gamma : \sigma'(A') \leq A$. An analogous statement holds for computations.*

5.3 Constraint Solving

The second phase of our inference-and-elaboration algorithm is the constraint solver. It is defined by the `solve` function signature:

$$\text{solve}(\sigma; \mathcal{P}; \mathcal{Q}) = (\sigma', \mathcal{P}')$$

It takes three inputs: the substitution σ accumulated so far, a list of already processed constraints \mathcal{P} , and a queue of still to be processed constraints \mathcal{Q} . There are two outputs: the substitution σ' that solves the constraints and the residual constraints \mathcal{P}' . The substitutions σ and σ' contain four kinds of mappings: $\varsigma \mapsto \tau$, $\alpha \mapsto A$, $\delta \mapsto \Delta$ and $\omega \mapsto \gamma$ which instantiate respectively skeleton variables, type variables, dirt variables and coercion variables.

Theorem 5 (Correctness of Solving). *For any set \mathcal{Q} , the call `solve($\bullet; \bullet; \mathcal{Q}$)` either results in a failure, in which case \mathcal{Q} has no solutions, or returns (σ, \mathcal{P}) such that for any $\sigma' \models \mathcal{Q}$, there exists $\sigma'' \models \mathcal{P}$ such that $\sigma' = \sigma'' \cdot \sigma$.*

The solver is invoked with `solve($\bullet; \bullet; \mathcal{Q}$)`, to process the constraints \mathcal{Q} generated in the first phase of the algorithm, i.e., with an empty substitution and no processed constraints. The `solve` function is defined by case analysis on the queue.

Empty Queue. When the queue is empty, all constraints have been processed. What remains are the residual constraints and the solving substitution σ , which are both returned as the result of the solver.

$$\text{solve}(\sigma; \mathcal{P}; \bullet) = (\sigma, \mathcal{P})$$

Skeleton Equalities. The next set of cases we consider are those where the queue is non-empty and its first element is an equality between skeletons $\tau_1 = \tau_2$. We consider seven possible cases based on the structure of τ_1 and τ_2 that together essentially implement conventional unification as used in Hindley-Milner type inference [5].

```

solve( $\sigma; \mathcal{P}; \tau_1 = \tau_2, \mathcal{Q}$ ) =
  match  $\tau_1 = \tau_2$  with
  |  $\varsigma = \varsigma \mapsto \text{solve}(\sigma; \mathcal{P}; \mathcal{Q})$ 
  |  $\varsigma = \tau \mapsto \text{if } \varsigma \notin \text{fv}_\varsigma(\tau) \text{ then let } \sigma' = [\tau/\varsigma] \text{ in } \text{solve}(\sigma' \cdot \sigma; \bullet; \sigma'(\mathcal{Q}, \mathcal{P})) \text{ else fail}$ 
  |  $\tau = \varsigma \mapsto \text{if } \varsigma \notin \text{fv}_\varsigma(\tau) \text{ then let } \sigma' = [\tau/\varsigma] \text{ in } \text{solve}(\sigma' \cdot \sigma; \bullet; \sigma'(\mathcal{Q}, \mathcal{P})) \text{ else fail}$ 
  |  $\text{Unit} = \text{Unit} \mapsto \text{solve}(\sigma; \mathcal{P}; \mathcal{Q})$ 
  |  $(\tau_1 \rightarrow \tau_2) = (\tau_3 \rightarrow \tau_4) \mapsto \text{solve}(\sigma; \mathcal{P}; \tau_1 = \tau_3, \tau_2 = \tau_4, \mathcal{Q})$ 
  |  $(\tau_1 \Rightarrow \tau_2) = (\tau_3 \Rightarrow \tau_4) \mapsto \text{solve}(\sigma; \mathcal{P}; \tau_1 = \tau_3, \tau_2 = \tau_4, \mathcal{Q})$ 
  | otherwise  $\mapsto \text{fail}$ 

```

The first case applies when both skeletons are the same type variable ς . Then the equality trivially holds. Hence we drop it and proceed with solving the remaining constraints. The next two cases apply when either τ_1 or τ_2 is a skeleton variable ς . If the occurs check fails, there is no finite solution and the algorithm signals failure. Otherwise, the constraint is solved by instantiating the ς . This additional substitution is accumulated and applied to all other constraints \mathcal{P}, \mathcal{Q} . Because the substitution might have modified some of the already processed constraints \mathcal{P} , we have to revisit them. Hence, they are all pushed back onto the queue, which is processed recursively.

The next three cases consider three different ways in which the two skeletons can have the same instantiated top-level structure. In those cases the equality is decomposed into equalities on the subterms, which are pushed onto the queue and processed recursively.

The last catch-all case deals with all ways in which the two skeletons can be instantiated to different structures. Then there is no solution.

Skeleton Annotations. The next four cases consider a skeleton annotation $\alpha : \tau$ at the head of the queue, and propagate the skeleton instantiation to the type variable. The first case, where the skeleton is a variable ς , has nothing to do, moves the annotation to the processed constraints and proceeds with the remainder of the queue. In the other three cases, the skeleton is instantiated and the solver instantiates the type variable with the corresponding structure, introducing fresh variables for any subterms. The instantiating substitution is accumulated and applied to the remaining constraints, which are processed recursively.

```

solve( $\sigma$ ;  $\mathcal{P}$ ;  $\alpha : \tau$ ,  $\mathcal{Q}$ ) =
  match  $\tau$  with
  |  $\varsigma$   $\mapsto$  solve( $\sigma$ ;  $\mathcal{P}$ ;  $\alpha : \tau$ ;  $\mathcal{Q}$ )
  | Unit  $\mapsto$  let  $\sigma' = [\text{Unit}/\alpha]$  in solve( $\sigma' \cdot \sigma$ ;  $\bullet$ ;  $\sigma'(\mathcal{Q}, \mathcal{P})$ )
  |  $\tau_1 \rightarrow \tau_2 \mapsto$  let  $\sigma' = [(\alpha_1^{\tau_1} \rightarrow \alpha_2^{\tau_2} ! \delta)/\alpha]$  in solve( $\sigma' \cdot \sigma$ ;  $\bullet$ ;  $\alpha_1 : \tau_1, \alpha_2 : \tau_2, \sigma'(\mathcal{Q}, \mathcal{P})$ )
  |  $\tau_1 \Rightarrow \tau_2 \mapsto$  let  $\sigma' = [(\alpha_1^{\tau_1} ! \delta_1 \Rightarrow \alpha_2^{\tau_2} ! \delta_2)/\alpha]$  in solve( $\sigma' \cdot \sigma$ ;  $\bullet$ ;  $\alpha_1 : \tau_1, \alpha_2 : \tau_2, \sigma'(\mathcal{Q}, \mathcal{P})$ )

```

Value Type Subtyping. Next are the cases where a subtyping constraint between two value types $A_1 \leq A_2$, with as evidence the coercion variable ω , is at the head of the queue. We consider six different situations.

```

solve( $\sigma$ ;  $\mathcal{P}$ ;  $\omega : A_1 \leq A_2$ ,  $\mathcal{Q}$ ) =
  match  $A_1 \leq A_2$  with
  |  $A \leq A \mapsto$  let  $T = \text{elab}_s(A)$  in solve( $[\langle T \rangle/\omega] \cdot \sigma$ ;  $\mathcal{P}$ ;  $\mathcal{Q}$ )
  |  $\alpha^{\tau_1} \leq A \mapsto$  let  $\tau_2 = \text{skeleton}(A)$  in solve( $\sigma$ ;  $\mathcal{P}$ ;  $\omega : \alpha^{\tau_1} \leq A$ ;  $\tau_1 = \tau_2$ ,  $\mathcal{Q}$ )
  |  $A \leq \alpha^{\tau_1} \mapsto$  let  $\tau_2 = \text{skeleton}(A)$  in solve( $\sigma$ ;  $\mathcal{P}$ ;  $\omega : A \leq \alpha^{\tau_1}$ ;  $\tau_2 = \tau_1$ ,  $\mathcal{Q}$ )
  |  $(A_1 \rightarrow B_1 ! \Delta_1) \leq (A_2 \rightarrow B_2 ! \Delta_2) \mapsto$  let  $\sigma' = [(\omega_1 \rightarrow \omega_2 ! \omega_3)/\omega]$  in
    solve( $\sigma' \cdot \sigma$ ;  $\mathcal{P}$ ;  $\omega_1 : A_2 \leq A_1, \omega_2 : B_1 \leq B_2, \omega_3 : \Delta_1 \leq \Delta_2$ ,  $\mathcal{Q}$ )
  |  $(A_1 ! \Delta_1 \Rightarrow A_2 ! \Delta_2) \leq (A_3 ! \Delta_3 \Rightarrow A_4 ! \Delta_4) \mapsto$  let  $\sigma' = [(\omega_1 ! \omega_2 \Rightarrow \omega_3 ! \omega_4)/\omega]$  in
    solve( $\sigma' \cdot \sigma$ ;  $\mathcal{P}$ ;  $\omega_1 : A_3 \leq A_1, \omega_2 : \Delta_3 \leq \Delta_1, \omega_3 : A_2 \leq A_4, \omega_4 : \Delta_2 \leq \Delta_4$ ,  $\mathcal{Q}$ )
  | otherwise  $\mapsto$  fail

```

If the two types are equal, the subtyping holds trivially through reflexivity. The solver thus drops the constraint and instantiates ω with the reflexivity coercion $\langle T \rangle$. Note that each coercion variable only appears in one constraint. So we only accumulate the substitution and do not have to apply it to the other constraints. In the next two cases, one of the two types is a type variable α . Then we move the constraint to the processed set. We also add an equality constraint between the skeletons⁸ to the queue. This enforces the invariant that only types with the same skeleton are compared. Through the skeleton equality the type structure (if any) from the type is also transferred to the type variable. The next two cases concern two types with the same top-level instantiation. The solver then decomposes the constraint into constraints on the corresponding subterms and appropriately relates the evidence of the old constraint to the new ones. The final case catches all situations where the two types are instantiated with a different structure and thus there is no solution.

Auxiliary function $\text{skeleton}(A)$ computes the skeleton of A .

Dirt Subtyping. The final six cases deal with subtyping constraints between dirts.

⁸ We implicitly annotate every type variable with its skeleton: α^τ .

```

solve( $\sigma$ ;  $\mathcal{P}$ ;  $\omega : \Delta \leq \Delta'$ ,  $\mathcal{Q}$ ) =
  match  $\Delta \leq \Delta'$  with
  |  $\mathcal{O} \cup \delta \leq \mathcal{O}' \cup \delta' \mapsto$  if  $\mathcal{O} \neq \emptyset$  then let  $\sigma' = [((\mathcal{O} \setminus \mathcal{O}') \cup \delta'')/\delta', \mathcal{O} \cup \omega'/\omega]$  in
      solve( $\sigma' \cdot \sigma$ ;  $\bullet$ ; ( $\omega' : \delta \leq \sigma'(\Delta')$ ),  $\sigma'(\mathcal{Q}, \mathcal{P})$ )
      else solve( $\sigma$ ;  $\mathcal{P}$ , ( $\omega : \Delta \leq \Delta'$ );  $\mathcal{Q}$ )
  |  $\emptyset \leq \Delta' \mapsto$  solve( $[\emptyset_{\Delta'}/\omega] \cdot \sigma$ ;  $\mathcal{P}$ ;  $\mathcal{Q}$ )
  |  $\delta \leq \emptyset \mapsto$  let  $\sigma' = [\emptyset/\delta; \emptyset_{\emptyset}/\omega]$  in solve( $\sigma' \cdot \sigma$ ;  $\bullet$ ;  $\sigma'(\mathcal{Q}, \mathcal{P})$ )
  |  $\mathcal{O} \cup \delta \leq \mathcal{O}' \mapsto$ 
      if  $\mathcal{O} \subseteq \mathcal{O}'$  then let  $\sigma' = [\mathcal{O} \cup \omega'/\omega]$  in solve( $\sigma' \cdot \sigma$ ;  $\mathcal{P}$ , ( $\omega' : \delta \leq \mathcal{O}'$ );  $\mathcal{Q}$ ) else fail
  |  $\mathcal{O} \leq \mathcal{O}' \mapsto$  if  $\mathcal{O} \subseteq \mathcal{O}'$  then let  $\sigma' = [\mathcal{O} \cup \emptyset_{\mathcal{O}' \setminus \mathcal{O}}/\omega]$  in solve( $\sigma' \cdot \sigma$ ;  $\mathcal{P}$ ;  $\mathcal{Q}$ ) else fail
  |  $\mathcal{O} \leq \mathcal{O}' \cup \delta' \mapsto$  let  $\sigma' = [(\mathcal{O} \setminus \mathcal{O}') \cup \delta''/\delta'; \mathcal{O}' \cup \emptyset_{(\mathcal{O}' \setminus \mathcal{O}) \cup \delta''}/\omega]$  in
      solve( $\sigma' \cdot \sigma$ ;  $\bullet$ ;  $\sigma'(\mathcal{Q}, \mathcal{P})$ )

```

If the two dirt sets are of the general form $\mathcal{O} \cup \delta$ and $\mathcal{O}' \cup \delta'$, we distinguish two subcases. Firstly, if \mathcal{O} is empty, there is nothing to be done and we move the constraint to the processed set. Secondly, if \mathcal{O} is non-empty, we partially instantiate δ' with any of the operations that appear in \mathcal{O} but not in \mathcal{O}' . We then drop \mathcal{O} from the constraint, and, after substitution, proceed with processing all constraints. For instance, for $\{\mathbf{Op}_1\} \cup \delta \leq \{\mathbf{Op}_2\} \cup \delta'$, we instantiate δ' to $\{\mathbf{Op}_1\} \cup \delta''$ —where δ'' is a fresh dirt variable—and proceed with the simplified constraint $\delta \leq \{\mathbf{Op}_1, \mathbf{Op}_2\} \cup \delta''$. Note that due to the set semantics of dirt sets, it is not valid to simplify the above constraint to $\delta \leq \{\mathbf{Op}_2\} \cup \delta''$. After all the substitution $[\delta \mapsto \{\mathbf{Op}_1\}, \delta'' \mapsto \emptyset]$ solves the former and the original constraint, but not the latter.

The second case, $\emptyset \leq \Delta'$, always holds and is discharged by instantiating ω to $\emptyset_{\Delta'}$. The third case, $\delta \leq \emptyset$, has only one solution: $\delta \mapsto \emptyset$ with coercion \emptyset_{\emptyset} . The fourth case, $\mathcal{O} \cup \delta \leq \mathcal{O}'$, has as many solutions as there are subsets of \mathcal{O}' , provided that $\mathcal{O} \subseteq \mathcal{O}'$. We then simplify the constraint to $\delta \leq \mathcal{O}'$, which we move to the set of processed constraints. The fifth case, $\mathcal{O} \leq \mathcal{O}'$, holds iff $\mathcal{O} \subseteq \mathcal{O}'$. The last case, $\mathcal{O} \leq \mathcal{O}' \cup \delta'$, is like the first, but without a dirt variable in the left-hand side. We can satisfy it in a similar fashion, by partially instantiating δ' with $(\mathcal{O} \setminus \mathcal{O}') \cup \delta''$ —where δ'' is a fresh dirt variable. Now the constraint is satisfied and can be discarded.

Terms

```

value  $v ::= x \mid \mathbf{unit} \mid h \mid \mathbf{fun} (x : \tau) \mapsto c \mid \Lambda_{\zeta}.v \mid v \tau$ 
handler  $h ::= \{\mathbf{return} (x : \tau) \mapsto c_r, \mathbf{Op}_1 x k \mapsto c_{\mathbf{Op}_1}, \dots, \mathbf{Op}_n x k \mapsto c_{\mathbf{Op}_n}\}$ 
computation  $c ::= v_1 v_2 \mid \mathbf{let} x = v \mathbf{in} c \mid \mathbf{return} v \mid \mathbf{Op} v (y : \tau.c)$ 
  |  $\mathbf{do} x \leftarrow c_1; c_2 \mid \mathbf{handle} c \mathbf{with} v$ 

```

Types

```

type  $\tau ::= \zeta \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \Rightarrow \tau_2 \mid \mathbf{Unit} \mid \forall_{\zeta}.\tau$ 

```

Fig. 10. SKELEff Syntax

5.4 Discussion

At first glance, the constraint generation algorithm of Sect. 5.2 might seem needlessly complex, due to eager constraint solving for let-generalization. Yet, we want to generalize at local `let`-bound values over both type and skeleton variables,⁹ which means that we must solve all equations between skeletons before generalizing. In turn, since skeleton constraints are generated when solving subtyping constraints (Sect. 5.3), all skeleton annotations should be available during constraint solving. This can not be achieved unless the generated constraints are propagated statefully.

6 Erasure of Effect Information from ExEff

6.1 The SkelEff Language

The target of the erasure is SKELEFF, which is essentially a copy of EXEFF from which all effect information Δ , type information T and coercions γ have been removed. Instead, skeletons τ play the role of plain types. Thus, SKELEFF is essentially System F extended with term-level (but not type-level) support for algebraic effects. Figure 10 defines the syntax of SKELEFF. The type system and operational semantics of SKELEFF follow from those of EXEFF.

Discussion. The main point of SKELEFF is to show that we can erase the effects and subtyping from EXEFF to obtain types that are compatible with a System F-like language. At the term-level SKELEFF also resembles a subset of Multicore OCaml [6], which provides native support for algebraic effects and handlers but features no explicit polymorphism. Moreover, SKELEFF can also serve as a staging area for further elaboration into System F-like languages without support for algebraic effects and handlers (e.g., Haskell or regular OCaml). In those cases, computation terms can be compiled to one of the known encodings in the literature, such as a free monad representation [10, 22], with delimited control [11], or using continuation-passing style [13], while values can typically be carried over as they are.

6.2 Erasure

Figure 11 defines erasure functions $\epsilon_v^\sigma(v)$, $\epsilon_c^\sigma(c)$, $\epsilon_V^\sigma(T)$, $\epsilon_C^\sigma(\underline{C})$ and $\epsilon_E^\sigma(\Gamma)$ for values, computations, value types, computation types, and type environments respectively. All five functions take a substitution σ from the free type variables α to their skeleton τ as an additional parameter.

Thanks to the skeleton-based design of EXEFF, erasure is straightforward. All types are erased to their skeletons, dropping quantifiers for type variables and all occurrences of dirt sets. Moreover, coercions are dropped from values

⁹ As will become apparent in Sect. 6, if we only generalize at the top over skeleton variables, the erasure does not yield local polymorphism.

$$\begin{array}{ll}
\epsilon_V^\sigma(x) = x & \epsilon_V^\sigma(\Lambda\delta.v) = \epsilon_V^\sigma(v) \\
\epsilon_V^\sigma(\mathbf{unit}) = \mathbf{unit} & \epsilon_V^\sigma(\Lambda(\omega : \pi).v) = \epsilon_V^\sigma(v) \\
\epsilon_V^\sigma(v \triangleright \gamma) = \epsilon_V^\sigma(v) & \epsilon_V^\sigma(v \tau) = \epsilon_V^\sigma(v) \tau \\
\epsilon_V^\sigma(\mathbf{fun} (x : T) \mapsto c) = \mathbf{fun} (x : \epsilon_V^\sigma(T)) \mapsto \epsilon_C^\sigma(c) & \epsilon_V^\sigma(v T) = \epsilon_V^\sigma(v) \\
\epsilon_V^\sigma(\Lambda\varsigma.v) = \Lambda\varsigma.\epsilon_V^\sigma(v) & \epsilon_V^\sigma(v \Delta) = \epsilon_V^\sigma(v) \\
\epsilon_V^\sigma(\Lambda(\alpha : \tau).v) = \epsilon_V^{\sigma.\{\alpha \mapsto \tau\}}(v) & \epsilon_V^\sigma(v \gamma) = \epsilon_V^\sigma(v) \\
\\
\epsilon_V^\sigma(\{\mathbf{return} (x : T) \mapsto c_r, [\mathbf{Op} x k \mapsto c_{\mathbf{Op}}]_{\mathbf{Op} \in \mathcal{O}}\}) = \\
\{\mathbf{return} (x : \epsilon_V^\sigma(T)) \mapsto \epsilon_C^\sigma(c_r), [\mathbf{Op} x k \mapsto \epsilon_C^\sigma(c_{\mathbf{Op}})]_{\mathbf{Op} \in \mathcal{O}}\} \\
\\
\epsilon_C^\sigma(v_1 v_2) = \epsilon_V^\sigma(v_1) \epsilon_V^\sigma(v_2) \\
\epsilon_C^\sigma(\mathbf{let} x = v \mathbf{in} c) = \mathbf{let} x = \epsilon_V^\sigma(v) \mathbf{in} \epsilon_C^\sigma(c) \\
\epsilon_C^\sigma(\mathbf{return} v) = \mathbf{return} (\epsilon_V^\sigma(v)) \\
\epsilon_C^\sigma(\mathbf{Op} v (y : T.c)) = \mathbf{Op} (\epsilon_V^\sigma(v)) (y : \epsilon_V^\sigma(T).\epsilon_C^\sigma(c)) \\
\epsilon_C^\sigma(\mathbf{do} x \leftarrow c_1; c_2) = \mathbf{do} x \leftarrow \epsilon_C^\sigma(c_1); \epsilon_C^\sigma(c_2) \\
\epsilon_C^\sigma(\mathbf{handle} c \mathbf{with} v) = \mathbf{handle} \epsilon_C^\sigma(c) \mathbf{with} \epsilon_V^\sigma(v) \\
\epsilon_C^\sigma(c \triangleright \gamma) = \epsilon_C^\sigma(c) \\
\\
\epsilon_V^\sigma(\alpha) = \sigma(\alpha) & \epsilon_C^\sigma(T ! \Delta) = \epsilon_V^\sigma(T) \\
\epsilon_V^\sigma(T \rightarrow \mathcal{C}) = \epsilon_V^\sigma(T) \rightarrow \epsilon_C^\sigma(\mathcal{C}) & \\
\epsilon_V^\sigma(\mathcal{C}_1 \cong \mathcal{C}_2) = \epsilon_C^\sigma(\mathcal{C}_1) \cong \epsilon_C^\sigma(\mathcal{C}_2) & \epsilon_E^\sigma(\epsilon) = \epsilon \\
\epsilon_V^\sigma(\mathbf{Unit}) = \mathbf{Unit} & \epsilon_E^\sigma(\Gamma, \varsigma) = \epsilon_E^\sigma(\Gamma), \varsigma \\
\epsilon_V^\sigma(\pi \Rightarrow T) = \epsilon_V^\sigma(T) & \epsilon_E^\sigma(\Gamma, \alpha : \tau) = \epsilon_E^{\sigma.\{\alpha \mapsto \tau\}}(\Gamma) \\
\epsilon_V^\sigma(\forall\varsigma.T) = \forall\varsigma.\epsilon_V^\sigma(T) & \epsilon_E^\sigma(\Gamma, \delta) = \epsilon_E^\sigma(\Gamma) \\
\epsilon_V^\sigma(\forall(\alpha : \tau).T) = \epsilon_V^{\sigma.\{\alpha \mapsto \tau\}}(T) & \epsilon_E^\sigma(\Gamma, x : T) = \epsilon_E^\sigma(\Gamma), x : \epsilon_V^\sigma(T) \\
\epsilon_V^\sigma(\forall\delta.T) = \epsilon_V^\sigma(T) & \epsilon_E^\sigma(\Gamma, \omega : \pi) = \epsilon_E^\sigma(\Gamma)
\end{array}$$

Fig. 11. Definition of type erasure.

and computations. Finally, all binders and elimination forms for type variables, dirt set variables and coercions are dropped from values and type environments.

The expected theorems hold. Firstly, types are preserved by erasure.¹⁰

Theorem 6 (Type Preservation). *If $\Gamma \vdash_v v : T$ then $\epsilon_E^\emptyset(\Gamma) \vdash_{\text{ev}} \epsilon_V^\Gamma(v) : \epsilon_V^\Gamma(T)$. If $\Gamma \vdash_c c : \underline{\mathcal{C}}$ then $\epsilon_E^\emptyset(\Gamma) \vdash_{\text{ec}} \epsilon_C^\Gamma(c) : \epsilon_C^\Gamma(\underline{\mathcal{C}})$.*

Here we abuse of notation and use Γ as a substitution from type variables to skeletons used by the erasure functions.

Finally, we have that erasure preserves the operational semantics.

Theorem 7 (Semantic Preservation). *If $v \rightsquigarrow_v v'$ then $\epsilon_V^\sigma(v) \equiv_v^{\rightsquigarrow} \epsilon_V^\sigma(v')$. If $c \rightsquigarrow_c c'$ then $\epsilon_C^\sigma(c) \equiv_c^{\rightsquigarrow} \epsilon_C^\sigma(c')$.*

In both cases, $\equiv^{\rightsquigarrow}$ denotes the congruence closure of the step relation in SKEL-EFF. The choice of substitution σ does not matter as types do not affect the behaviour.

¹⁰ Typing for SKEL-EFF values and computations take the form $\Gamma \vdash_{\text{ev}} v : \tau$ and $\Gamma \vdash_{\text{ec}} c : \tau$.

Discussion. Typically, when type information is erased from call-by-value languages, type binders are erased by replacing them with other (dummy) binders. For instance, the expected definition of erasure would be:

$$\epsilon_v^\sigma(\lambda(\alpha : \tau).v) = \lambda(x : \mathbf{Unit}).\epsilon_v^\sigma(v)$$

This replacement is motivated by a desire to preserve the behaviour of the typed terms. By dropping binders, values might be turned into computations that trigger their side-effects immediately, rather than at the later point where the original binder was eliminated. However, there is no call for this circumspect approach in our setting, as our grammatical partition of terms in values (without side-effects) and computations (with side-effects) guarantees that this problem cannot happen when we erase values to values and computations to computations.

7 Related Work and Conclusion

Eff's Implicit Type System. The most closely related work is that of Pretnar [20] on inferring algebraic effects for Eff, which is the basis for our implicitly-typed IMPEFF calculus, its type system and the type inference algorithm. There are three major differences with Pretnar's inference algorithm.

Firstly, our work introduces an explicitly-typed calculus. For this reason we have extended the constraint generation phase with the elaboration into EXEFF and the constraint solving phase with the construction of coercions.

Secondly, we add skeletons to guarantee erasure. Skeletons also allow us to use standard occurs-check during unification. In contrast, unification in Pretnar's algorithm is inspired by Simonet [24] and performs the occurs-check up to the equivalence closure of the subtyping relation. In order to maintain invariants, all variables in an equivalence class (also called a skeleton) must be instantiated simultaneously, whereas we can process one constraint at a time. As these classes turn out to be surrogates for the underlying skeleton types, we have decided to keep the name.

Finally, Pretnar incorporates garbage collection of constraints [19]. The aim of this approach is to obtain unique and simple type schemes by eliminating redundant constraints. Garbage collection is not suitable for our use as type variables and coercions witnessing subtyping constraints cannot simply be dropped, but must be instantiated in a suitable manner, which cannot be done in general.

Consider for instance a situation with type variables $\alpha_1, \alpha_2, \alpha_3, \alpha_4$, and α_5 where $\alpha_1 \leq \alpha_3$, $\alpha_2 \leq \alpha_3$, $\alpha_3 \leq \alpha_4$, and $\alpha_3 \leq \alpha_5$. Suppose that α_3 does not appear in the type. Then garbage collection would eliminate it and replace the constraints by $\alpha_1 \leq \alpha_4$, $\alpha_2 \leq \alpha_4$, $\alpha_1 \leq \alpha_5$, and $\alpha_2 \leq \alpha_5$. While garbage collection guarantees that for any ground instantiation of the remaining type variables, there exists a valid ground instantiation for α_3 , EXEFF would need to be extended with joins (or meets) to express a generically valid instantiation like $\alpha_1 \sqcup \alpha_2$. Moreover, we would need additional coercion formers to establish $\alpha_1 \leq (\alpha_1 \sqcup \alpha_2)$ or $(\alpha_1 \sqcup \alpha_2) \leq \alpha_4$.

As these additional constructs considerably complicate the calculus, we propose a simpler solution. We use EXEFF as it is for internal purposes, but display types to programmers in their garbage-collected form.

Calculi with Explicit Coercions. The notion of explicit coercions is not new; Mitchell [15] introduced the idea of inserting coercions during type inference for ML-based languages, as a means for explicit casting between different numeric types.

Breazu-Tannen et al. [3] also present a translation of languages with inheritance polymorphism into System F, extended with coercions. Although their coercion combinators are very similar to our coercion forms, they do not include inversion forms, which are crucial for the proof of type safety for our system. Moreover, Breazu-Tannen et al.’s coercions are terms, and thus can not be erased.

Much closer to EXEFF is Crary’s coercion calculus for inclusive subtyping [4], from which we borrowed the stratification of value results. Crary’s system supports neither coercion abstraction nor coercion inversion forms.

System F_C [25] uses explicit type-equality coercions to encode complex language features (e.g. GADTs [16] or type families [23]). Though EXEFF’s coercions are proofs of subtyping rather than type equality, our system has a lot in common with it, including the inversion coercion forms and the “push” rules.

Future Work. Our plans focus on resuming the postponed work on efficient compilation of handlers. First, we intend to adjust program transformations to the explicit type information. We hope that this will not only make the optimizer more robust, but also expose new optimization opportunities. Next, we plan to write compilers to both Multicore OCaml and standard OCaml, though for the latter, we must first adapt the notion of erasure to a target calculus without algebraic effect handlers. Finally, once the compiler shows promising preliminary results, we plan to extend it to other Eff features such as user-defined types or recursion, allowing us to benchmark it on more realistic programs.

Acknowledgements. We would like to thank the anonymous reviewers for careful reading and insightful comments. Part of this work is funded by the Flemish Fund for Scientific Research (FWO). This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-17-1-0326.

References

1. Barendregt, H.: The Lambda Calculus: Its Syntax and Semantics. Studies in Logic and the Foundations of Mathematics, vol. 3. North-Holland, Amsterdam (1981)
2. Bauer, A., Pretnar, M.: Programming with algebraic effects and handlers. *J. Logic Algebraic Program.* **84**(1), 108–123 (2015)
3. Breazu-Tannen, V., Coquand, T., Gunter, C.A., Scedrov, A.: Inheritance as implicit coercion. *Inf. Comput.* **93**, 172–221 (1991)
4. Crary, K.: Typed compilation of inclusive subtyping. In: Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming, ICFP 2000, pp. 68–81. ACM, New York (2000)

5. Damas, L., Milner, R.: Principal type-schemes for functional programs. In: Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 1982, pp. 207–212. ACM, New York (1982)
6. Dolan, S., White, L., Sivaramakrishnan, K., Yallop, J., Madhavapeddy, A.: Effective concurrency through algebraic effects. In: OCaml Workshop (2015)
7. Girard, J.-Y., Taylor, P., Lafont, Y.: Proofs and Types. Cambridge University Press, Cambridge (1989)
8. Hillerström, D., Lindley, S.: Liberating effects with rows and handlers. In: Chapman, J., Swierstra, W. (eds.) Proceedings of the 1st International Workshop on Type-Driven Development, TyDe@ICFP 2016, Nara, Japan, 18 September 2016, pp. 15–27. ACM (2016)
9. Jones, M.P.: A theory of qualified types. In: Krieg-Brückner, B. (ed.) ESOP 1992. LNCS, vol. 582, pp. 287–306. Springer, Heidelberg (1992). <https://doi.org/10.1007/3-540-55253-7-17>
10. Kammar, O., Lindley, S., Oury, N.: Handlers in action. In: Proceedings of the 18th ACM SIGPLAN International Conference on Functional programming, ICFP 2014, pp. 145–158. ACM (2013)
11. Kiselyov, O., Sivaramakrishnan, K.: Eff directly in OCaml. In: OCaml Workshop (2016)
12. Leijen, D.: Koka: programming with row polymorphic effect types. In: Levy, P., Krishnaswami, N. (eds.) Proceedings of 5th Workshop on Mathematically Structured Functional Programming, MSFP@ETAPS 2014, Grenoble, France, 12 April 2014. EPTCS, vol. 153, pp. 100–126 (2014)
13. Leijen, D.: Type directed compilation of row-typed algebraic effects. In: Castagna, G., Gordon, A.D. (eds.) Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, 18–20 January 2017, pp. 486–499. ACM (2017)
14. Lindley, S., McBride, C., McLaughlin, C.: Do be do be do. In: Castagna, G., Gordon, A.D. (eds.) Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, 18–20 January 2017, pp. 500–514. ACM (2017). <http://dl.acm.org/citation.cfm?id=3009897>
15. Mitchell, J.C.: Coercion and type inference. In: Proceedings of the 11th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages, POPL 1984, pp. 175–185. ACM, New York (1984)
16. Peyton Jones, S., Vytiniotis, D., Weirich, S., Washburn, G.: Simple unification-based type inference for GADTs. In: ICFP 2006 (2006)
17. Plotkin, G.D., Power, J.: Algebraic operations and generic effects. Appl. Categ. Struct. **11**(1), 69–94 (2003)
18. Plotkin, G.D., Pretnar, M.: Handling algebraic effects. Log. Methods Comput. Sci. **9**(4) (2013)
19. Pottier, F.: Simplifying subtyping constraints: a theory. Inf. Comput. **170**(2), 153–183 (2001). <https://doi.org/10.1006/inco.2001.2963>
20. Pretnar, M.: Inferring algebraic effects. Log. Methods Comput. Sci. **10**(3) (2014)
21. Pretnar, M.: An introduction to algebraic effects and handlers, invited tutorial. Electron. Notes Theoret. Comput. Sci. **319**, 19–35 (2015)
22. Pretnar, M., Saleh, A.H., Faes, A., Schrijvers, T.: Efficient compilation of algebraic effects and handlers. Technical report CW 708, KU Leuven Department of Computer Science (2017)
23. Schrijvers, T., Peyton Jones, S., Chakravarty, M., Sulzmann, M.: Type checking with open type functions. In: ICFP 2008, pp. 51–62. ACM (2008)

24. Simonet, V.: Type inference with structural subtyping: a faithful formalization of an efficient constraint solver. In: Ohori, A. (ed.) APLAS 2003. LNCS, vol. 2895, pp. 283–302. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-40018-9_19
25. Sulzmann, M., Chakravarty, M.M.T., Peyton Jones, S., Donnelly, K.: System F with type equality coercions. In: Proceedings of the 2007 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation, TLDI 2007, pp. 53–66. ACM, New York (2007)
26. Wansbrough, K., Peyton Jones, S.L.: Once upon a polymorphic type. In: POPL, pp. 15–28. ACM (1999)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

