

Asymptotic Improvement with Effect Handlers

Daniel Hillerström
The University of Edinburgh
UK
daniel.hillerstrom@ed.ac.uk

Sam Lindley
The University of Edinburgh and
Imperial College London
UK
sam.lindley@ed.ac.uk

John Longley
The University of Edinburgh
UK
jrl@staffmail.ed.ac.uk

Abstract

As Filinski showed in the 1990s, delimited control operators can express all monadic effects. Plotkin and Pretnar’s effect handlers offer a modular form of delimited control providing a uniform mechanism for concisely implementing features ranging from `async/await` to probabilistic programming.

We study the fundamental efficiency of effect handlers. Specifically, we show that effect handlers enable an asymptotic improvement in runtime complexity for a certain class of programs. We consider the *generic search* problem and define a pure PCF-like base language λ_b and its extension with effect handlers λ_h . We show that λ_h admits an asymptotically more efficient implementation of generic search than any λ_b implementation of generic search.

To our knowledge this result is the first of its kind for control operators.

1 Introduction

In today’s programming languages we find a wealth of powerful constructs and features — exceptions, higher-order store, dynamic method dispatch, coroutining, explicit continuations, concurrency features, Lisp-style ‘quote’ and so on — which may be present or absent in various combinations in any given language. There are of course many important pragmatic and stylistic differences between languages, but here we are concerned with whether languages may differ more essentially in their expressive power, according to the selection of features they contain.

One can interpret this question in various ways. For instance, Felleisen [13] considers the question of whether a language \mathcal{L} admits a translation into a sublanguage \mathcal{L}' in a way which respects not only the behaviour of programs but also aspects of their (global or local) syntactic structure. If the translation of some \mathcal{L} -program into \mathcal{L}' requires a complete global restructuring, we may say that \mathcal{L}' is in some way less expressive than \mathcal{L} . In the present paper, however, we have in mind even more fundamental expressivity differences that would not be bridged even if whole-program translations were admitted. These fall under two headings.

1. *Computability*: Are there operations of type A that are programmable in \mathcal{L} but not expressible at all in \mathcal{L}' ?
2. *Complexity*: Are there operations programmable in \mathcal{L} with some asymptotic runtime bound (e.g. ‘ $O(n^2)$ ’) that cannot be achieved in \mathcal{L}' ?

We may also ask: are there examples of *natural, practically useful* operations that manifest such differences? If so, this might be considered as a significant advantage of \mathcal{L} over \mathcal{L}' .

If the ‘operations’ we are asking about are ordinary first-order functions — that is, both their inputs and outputs are of ground type (strings, arbitrary-size integers etc.) — then the situation is easily summarised. At such types, all reasonable languages give rise to the same class of programmable functions, namely the Church-Turing computable ones. As for complexity, the runtime of a program is typically analysed with respect to some cost model for basic instructions (e.g. one unit of time per array access). Although the realism of such cost models in the asymptotic limit can be questioned (see, e.g., [23, Section 2.6]), it is broadly taken as read that such models are equally applicable whatever programming language we are working with, and moreover that all respectable languages can represent all algorithms of interest; thus, one does not expect the best achievable asymptotic run-time for a typical algorithm (say in number theory or graph theory) to be sensitive to the choice of programming language, except perhaps in marginal cases.

The situation changes radically, however, if we consider *higher-order* operations: programmable operations whose inputs may themselves be programmable operations. Here it turns out that both what is computable and the efficiency with which it can be computed can be highly sensitive to the selection of language features present. This is in fact true more widely for *abstract data types*, of which higher-order types can be seen as a special case: a higher-order value will of course be represented within the machine as ground data, but a program within the language typically has no access to this internal representation, and can interact with the value only by applying it to an argument.

Most of the work in this area to date has focused on computability differences. One of the best known examples is the *parallel if* operation which is computable in a language with parallel evaluation but not in a typical ‘sequential’ programming language [35]. It is also well known that the presence of control features or local state enables observational distinctions that cannot be made in a purely functional setting: for instance, there are programs involving ‘call/cc’ that detect the order in which a (call-by-name) ‘+’ operation evaluates its arguments [9]. Such operations are ‘non-functional’ in the sense that their output is not determined solely by the extension of their input; however, there are also programs

with ‘functional’ behaviour that can be implemented with control or local state but not without them [28]. More recent results have exhibited differences lower down in the language expressivity spectrum: for instance, in a purely functional setting *à la* Haskell, the expressive power of *recursion* increases strictly with its type level [30], and there are natural operations computable by (low-order) recursion but not by (even high-order) iteration [29]. Much of this territory, including the mathematical theory of some of the natural notions of higher-order computability that arise in this way, is mapped out by Longley and Normann [31].

Relatively few results of this character have so far been established on the complexity side. Pippenger [33] gives an example of an ‘online’ operation on infinite sequences of atomic symbols (essentially a function from streams to streams) such that the first n output symbols can be produced within time $O(n)$ if one is working in an ‘impure’ version of Lisp (in which mutation of ‘cons’ pairs is admitted), but with a worst-case runtime no better than $\Omega(n \log n)$ for any implementation in pure Lisp (without such mutation). This example was reconsidered by Bird et al. [8] who showed that the same speedup can be achieved in a pure language by using lazy evaluation. Jones [20] explores the approach of manifesting expressivity and efficiency differences between certain languages by artificially restricting attention to ‘cons-free’ programs; in this setting, the classes of representable first-order functions for the various languages are found to coincide with some well-known complexity classes.

The purpose of the present paper is to give a clear example of such an inherent complexity difference higher up in the expressivity spectrum. Specifically, we consider the following *generic search* problem, parametric in n : given a boolean-valued predicate P on the space \mathbb{B}^n of boolean vectors of length n , return the number of such vectors p for which $P(p) = \text{true}$. We shall consider boolean vectors of any length to be represented by the type $\text{Nat} \rightarrow \text{Bool}$; thus, for each n , we are asking for an implementation of a certain third-order operation

$$\text{count}_n : ((\text{Nat}_n \rightarrow \text{Bool}) \rightarrow \text{Bool}) \rightarrow \text{Nat}$$

A naive implementation strategy, supported by any reasonable language, is simply to apply P to each of the 2^n vectors in turn. A much less obvious, but still purely ‘functional’, approach due to Berger [5] achieves the effect of ‘pruned search’ where the predicate admits this (serving as a warning that counter-intuitive phenomena can arise in this territory). Nonetheless, under a mild condition on P (namely that it must inspect all n components of the given vector before returning), both these approaches will have a $\Omega(n2^n)$ runtime. Moreover, we shall show that in a typical call-by-value language without advanced control features, one cannot improve on this: *any* implementation of count_n must necessarily take time $\Omega(n2^n)$, even when the predicates P are chosen to be ‘as simple as possible’. On the other hand, if

we extend our language with a feature such as *effect handlers* (see Section 2 below), it becomes possible to bring the runtime down to $O(2^n)$: an asymptotic gain of a factor of n .

In order to make this efficiency difference stand out as clearly as possible, we have resorted to some slightly artificial constraints in the way we set up our scenario. Of course, even one illustration of this phenomenon suffices in principle to establish the existence of the efficiency gap in question; however, it will also be clear from our analysis that, in spite of the technical restrictions, the phenomenon we are exhibiting is actually quite general. For instance, if the problem of counting all solutions to P is replaced by that of returning the first solution found (if one exists), it will be clear that an order n speedup can still typically be expected.

The idea behind the speedup is easily explained. Suppose for example $n = 4$, and suppose that the predicate P always inspects the components of its argument in the order 0, 1, 2, 3. A naive implementation of count_4 might start by applying the given P to $p_0 = (\text{true}, \text{true}, \text{true}, \text{true})$, and then to $p_1 = (\text{true}, \text{true}, \text{true}, \text{false})$. Clearly there is some duplication here: the computations of $P p_0$ and $P p_1$ will proceed identically up to the point where the value of the final component is requested. What we would like to do, then, is to record the state of the computation of $P p_0$ at just this point, so that we can later resume this computation with false supplied as the final component value in order to obtain the value of $P p_1$. (Similarly for all other internal nodes in the evident binary tree of boolean vectors.) Of course, this ‘backup’ approach would be standardly applied if one were implementing a bespoke search operation for some *particular* choice of P (corresponding, say, to the n -queens problem); but to apply this idea of resuming previous subcomputations in the generic setting (that is, uniformly in P) requires some special language feature such as effect handlers or multi-use continuations. One could also obviate the need for such a feature by choosing to present the predicate P in some other way, but from our present perspective this would be to move the goalposts: our intention is precisely to show that our languages differ in an essential way *as regards their power to manipulate data of type* $(\text{Nat} \rightarrow \text{Bool}) \rightarrow \text{Bool}$.

The above idea will already be familiar, at least informally, to many who have worked with effect handlers or explicit continuations, and was explicitly presented in a closely related context by Bauer [1]; but our contribution here is to formulate and prove a precise mathematical theorem that pins down the efficiency difference in question. A general mathematical theory of the expressive power of effect handlers would be perhaps best articulated within the framework of game semantics; since in the present paper our focus is on one specific example of the difference, we shall work concretely and operationally with the languages themselves. In the first instance, we formulate our results as a comparison between a purely functional base language (a version of call-by-value PCF [35]) and an extension of this with effect

handlers; we then easily observe that our results are unaffected if the base language is augmented with other features such as local mutable store. As regards the runtime estimates, we work with a CEK-style abstract machine model for our languages which, we claim, offers a realistic model of program execution time for typical real-world implementations.

The rest of the paper is structured as follows.

- Section 2 provides an introduction to effect handlers as a programming abstraction.
- Section 3 presents a PCF-like language λ_b and its extension λ_h with effect handlers.
- Section 4 defines abstract machines for λ_b and λ_h .
- Section 5 formally states and proves the complexity of generic search in λ_b ($\Omega(n2^n)$) and λ_h ($O(2^n)$).
- Section 6 outlines how the result scales to extensions of the base language with features such as state.
- Section 7 empirically evaluates implementations of generic search based on λ_b and λ_h in Standard ML.
- Section 8 concludes.

The languages λ_b and λ_h are rather minimal variants of previous work – we only include the machinery needed for illustrating the generic search efficiency phenomenon. Full proofs of our main complexity results are available in the appendices of the anonymised supplementary material.

2 Effect Handlers Primer

Effect handlers were originally studied as a theoretical means to provide a semantics for exception handling in the setting of algebraic effects [36, 37]. Subsequently they have emerged as a practical programming abstraction for modular effectful programming [3, 11, 18, 21, 22, 24, 27]. In this section we give a short introduction to effect handlers. For a thorough introduction to programming with effect handlers, we recommend the tutorial by Pretnar [38], and as an introduction to the mathematical foundations of handlers, we refer the reader to the founding paper by Plotkin and Pretnar [37] and the excellent tutorial paper by Bauer [2].

Viewed through the lens of universal algebra, an algebraic effect is given by a signature Σ of finitary *operation symbols* defined over some nonempty carrier set A , along with an equational theory that describes the properties of the operations [36]. An example of an algebraic effect is *nondeterminism*, whose signature consists of a single nondeterministic choice operation: $\Sigma := \{\text{Branch} : 1 \rightarrow \text{Bool}\}$. The operation takes a single parameter of type unit and ultimately produces a boolean value. The pragmatic programmatic view of algebraic effects differs from the original development as no implementation accounts for equations over operations yet.


As an introductory programmatic example, we will showcase a use of the operation `Branch` by modelling a coin toss. Suppose we have a data type $\text{Toss} := \text{Heads} \mid \text{Tails}$, then

in our programming notation (introduced formally in Section 3.2) we may implement a coin toss as follows.

```
toss : ⟨ ⟩ → Toss
toss ⟨ ⟩ = if do Branch ⟨ ⟩ then inl Heads else inr Tails
```

From the type signature it is clear that the computation returns a value of type `Toss`. It is not clear from the signature of `toss` whether it performs an effect. From looking at the definition, it evidently performs the operation `Branch` with argument $\langle \rangle$ using the `do`-invocation form. The result of the operation determines whether the computation returns either `Heads` or `Tails` (with the appropriate injections). Systems such as Frank [27], Helium [7], Koka [24], and Links [18] include type-and-effect systems which track the use of effectful operations. Whilst current iterations of systems such as Eff [3] and Multicore OCaml [11] elect not to include an effect system. Our language is closer to the latter two.

We may, in the style of Lindley [26], view an effectful computation as a tree, where the interior nodes correspond to operation invocations and the leaves correspond to return values. The computation tree for `toss` is as follows.



It models interaction with the environment. The operation `Branch` can be viewed as a *query* for which the *response* is either `true` or `false`. The response is provided by an effect handler. As an example consider the following handler which enumerates the possible outcomes of a coin toss.

```
handle toss ⟨ ⟩ with
  val x      ↦ [x]
  Branch ⟨ ⟩ r ↦ r true ++ r false
```

The `handle`-construct generalises the exceptional syntax of Benton and Kennedy [4]. A handler has a *success* clause and an *operation* clause. The success clause determines how to interpret the return value of `toss`. It lifts the return value into a singleton list. The operation clause determines how to interpret occurrences of `Branch` in `toss`. It provides access to the argument of `Branch` (which is unit) and its resumption, r . The resumption is a first-class delimited continuation which captures the remainder of the `toss` computation from the invocation of `Branch` up to its nearest enclosing handler.

Applying r to `true` resumes evaluation of `toss` via the true branch, returning `Heads` and causing the success clause of the handler to be invoked; thus the result of $r \text{ true}$ is `[inl Heads]`. Evaluation continues in the operation clause, meaning that r is applied again, but this time to `false`, which causes evaluation to resume in `toss` via the false branch. By the same reasoning, the value of $r \text{ false}$ is `[inr Tails]`, which is concatenated with the result of the true branch; hence the handler ultimately returns `[inl Heads, inr Tails]`.

3 Calculi

In this section, we present our base language λ_b and its extension with effect handlers λ_h .

3.1 Base Calculus

The base calculus λ_b is a fine-grain call-by-value [25] variation of PCF [35]. Fine-grain call-by-value is similar to A-normal form [16] in that every intermediate computation is named, but unlike A-normal form is closed under reduction.

The types of λ_b are given by the following grammar.

$$A, B, C, D ::= \text{Nat} \mid 1 \mid A \rightarrow B \mid A \times B \mid A + B$$

The ground types are Nat and 1 which classify natural number values and the unit value, respectively. We write ground A to assert that type A is a ground type. The function type $A \rightarrow B$ represents functions that map values of type A to values of type B . The binary product type $A \times B$ represents a pair of values whose first and second components have types A and B respectively. The sum type $A + B$ represents tagged values of either type A or B . Type environments Γ map term variables to their types.

We let n range over natural numbers and c range over primitive operations on natural numbers ($+$, $-$, $=$). We generally use lowercase letters x, y, z and more to denote term variables. By convention we use $f, g,$ and h for variables of function type, i and j for variables of type Nat , and r and k to denote resumptions and continuations, with the exception that we will use uppercase P to denote predicates.

The typing rules are given in Figure 1. We require two typing judgements: one for values and the other for computations. The judgement $\Gamma \vdash \square : A$ states that a \square -term has type A under type environment Γ , where \square is either a value term (V) or a computation term (M). The constants have the following types.

$$\{(+), (-)\} : \langle \text{Nat}, \text{Nat} \rangle \rightarrow \text{Nat} \quad (=) : \langle \text{Nat}, \text{Nat} \rangle \rightarrow \text{Bool}$$

Value terms comprise variables (x), the unit value ($\langle \rangle$), natural number literals (n), primitive constants (c), lambda abstraction ($\lambda x^A. M$), recursion ($\mathbf{rec} f^A x. M$), pairs ($\langle V, W \rangle$), left ($\mathbf{inl} V^B$) and right ($\mathbf{inr} W^A$) injections. We assume an efficient representation of naturals (e.g. naturals occupy a machine word) such that constants ($c \in \{+, -, =\}$) have efficient realisations $\ulcorner c \urcorner$. All elimination forms are computation terms. Abstraction is eliminated using application ($V W$). The product eliminator ($\mathbf{let} \langle x, y \rangle = V \mathbf{in} N$) splits a pair V into its constituents and binds them to x and y , respectively. Sums are eliminated by a case split ($\mathbf{case} V \{\mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N\}$). A trivial computation ($\mathbf{return} V$) returns value V . The sequencing expression ($\mathbf{let} x \leftarrow M \mathbf{in} N$) evaluates M and binds the result value to x in N .

For convenience we often write code in direct-style assuming the standard left-to-right call-by-value elaboration into fine-grain call-by-value [16]. For example, assuming f, g, h are functions, and a is a bound variable, then the expression

$(f(h a) + g \langle \rangle)$ is syntactic sugar for:

$$\mathbf{let} x \leftarrow h a \mathbf{in} \mathbf{let} y \leftarrow f x \mathbf{in} \mathbf{let} z \leftarrow g \langle \rangle \mathbf{in} y + z$$

We use the standard encoding of booleans as sums:

$$\text{Bool} := 1 + 1 \quad \text{true} := \mathbf{inl} \langle \rangle \quad \text{false} := \mathbf{inr} \langle \rangle$$

$$\mathbf{if} V \mathbf{then} M \mathbf{else} N := \mathbf{case} V \{\mathbf{inl} \langle \rangle \mapsto M; \mathbf{inr} \langle \rangle \mapsto N\}$$

We make use of standard syntactic sugar for pattern matching. For instance, for suspended computations we write

$$\lambda \langle \rangle. M := \lambda x^1. M, \quad \text{where } x \notin FV(M)$$

and more generally if the binder has a type other than 1 , then we write

$$\lambda_{-}^A. M := \lambda x^A. M, \quad \text{where } x \notin FV(M)$$

We elide type annotations when clear from context.

We give a small-step operational semantics for λ_b with *evaluation contexts* in the style of Felleisen [12].

$$(\lambda x^A. M) V \rightsquigarrow M[V/x]$$

$$(\mathbf{rec} f^A x. M) V \rightsquigarrow M[(\mathbf{rec} f^A x. M)/f, V/x]$$

$$c V \rightsquigarrow \mathbf{return} (\ulcorner c \urcorner(V))$$

$$\mathbf{let} \langle x; y \rangle = \langle V; W \rangle \mathbf{in} N \rightsquigarrow N[V/x, W/y]$$

$$\mathbf{case} (\mathbf{inl} V)^B \{\mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N\} \rightsquigarrow M[V/x]$$

$$\mathbf{case} (\mathbf{inr} V)^A \{\mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N\} \rightsquigarrow N[V/y]$$

$$\mathbf{let} x \leftarrow \mathbf{return} V \mathbf{in} N \rightsquigarrow N[V/x]$$

$$\mathcal{E}[M] \rightsquigarrow \mathcal{E}[N], \quad \text{if } M \rightsquigarrow N$$

$$\text{Evaluation contexts } \mathcal{E} ::= [] \mid \mathbf{let} x \leftarrow \mathcal{E} \mathbf{in} N$$

We write $M[V/x]$ for M with V substituted for x . The reduction relation \rightsquigarrow is defined on computation terms. The statement $M \rightsquigarrow N$ reads: term M reduces to term N in one step. We write R^+ for the transitive closure of relation R .

3.2 Handler Calculus

We now define λ_h as an extension of λ_b . First we define notation for operation symbols, signatures, and handler types.

$$\text{Operation symbols } \ell \in \mathcal{L}$$

$$\text{Signatures } \Sigma ::= \cdot \mid \{\ell : A \rightarrow B\} \cup \Sigma$$

$$\text{Handler types } F ::= C \Rightarrow D$$

We assume a countably infinite set of operation symbols \mathcal{L} . An effect signature Σ is a map from operation symbols to their types, thus we assume that each operation symbol in a signature is distinct. An operation type $A \rightarrow B$ denotes an operation that takes an argument of type A and returns a result of type B . We write $\text{dom}(\Sigma) \subseteq \mathcal{L}$ for the set of operation symbols in a signature Σ . An effect handler type $C \Rightarrow D$ classifies effect handlers that transform computations of type C into computations of type D . Following Pretnar [38], we assume a global signature for every program.

The typing rules for λ_h are those of λ_b (Figure 1) plus three additional rules for operations, handling, and handlers given in Figure 2. The T-OP rule ensures that an operation invocation is only well-typed if the operation ℓ appears in the

441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495

Values

$\frac{\Gamma \vdash V : A}{\Gamma \vdash x : A}$	$\frac{\text{T-UNIT}}{\Gamma \vdash \langle \rangle : 1}$	$\frac{\text{T-NAT} \quad n \in \mathbb{N}}{\Gamma \vdash n : \text{Nat}}$	$\frac{\text{T-CONST} \quad c : A \rightarrow B}{\Gamma \vdash c : A \rightarrow B}$
$\frac{\text{T-LAM} \quad \Gamma, x : A \vdash M : C}{\Gamma \vdash \lambda x^A. M : A \rightarrow C}$	$\frac{\text{T-REC} \quad \Gamma, f : A \rightarrow C, x : A \vdash M : C}{\Gamma \vdash \text{rec } f^{A \rightarrow C} x. M : A \rightarrow C}$		
$\frac{\text{T-PROD} \quad \Gamma \vdash V : A \quad \Gamma \vdash W : B}{\Gamma \vdash \langle V, W \rangle : A \times B}$	$\frac{\text{T-INL} \quad \Gamma \vdash V : A}{\Gamma \vdash (\text{inl } V)^B : A + B}$	$\frac{\text{T-INR} \quad \Gamma \vdash W : B}{\Gamma \vdash (\text{inr } W)^A : A + B}$	

Computations

$\frac{\text{T-APP} \quad \Gamma \vdash V : A \rightarrow B \quad \Gamma \vdash W : A}{\Gamma \vdash V W : B}$	$\frac{\text{T-SPLIT} \quad \Gamma \vdash V : A \times B \quad \Gamma, x : A, y : B \vdash N : C}{\Gamma \vdash \text{let } \langle x, y \rangle = V \text{ in } N : C}$
$\frac{\text{T-CASE} \quad \Gamma \vdash V : A + B \quad \Gamma, x : A \vdash M : C \quad \Gamma, y : B \vdash N : C}{\Gamma \vdash \text{case } V \{ \text{inl } x \mapsto M; \text{inr } y \mapsto N \} : C}$	
$\frac{\text{T-RETURN} \quad \Gamma \vdash V : A}{\Gamma \vdash \text{return } V : A}$	$\frac{\text{T-LET} \quad \Gamma \vdash M : A \quad \Gamma, x : A \vdash N : C}{\Gamma \vdash \text{let } x \leftarrow M \text{ in } N : C}$

Figure 1. Typing Rules for λ_b **Computations**

$\frac{\text{T-DO} \quad (\ell : A \rightarrow B) \in \Sigma \quad \Gamma \vdash V : A}{\Gamma \vdash \text{do } \ell V : B}$	$\frac{\text{T-HANDLE} \quad \Gamma \vdash M : C \quad \Gamma \vdash H : C \Rightarrow D}{\Gamma \vdash \text{handle } M \text{ with } H : D}$
---	--

Handlers

$\frac{\text{T-HANDLER} \quad H^{\text{val}} = \{ \text{val } x \mapsto M \} \quad [H^\ell = \{ \ell p_\ell r_\ell \mapsto N_\ell \}]_{\ell \in \text{dom}(\Sigma)} \quad [\Gamma, p_\ell : A_\ell, r_\ell : B_\ell \rightarrow D \vdash N_\ell : D]_{(\ell : A_\ell \rightarrow B_\ell) \in \Sigma}}{\Gamma, x : C \vdash M : D}$
$\Gamma \vdash H : C \Rightarrow D$

Figure 2. Additional Typing Rules for λ_h

effect signature Σ and the argument type A matches the type of the provided argument V . The result type B determines the type of the invocation. The T-HANDLE rule is straightforward. The T-HANDLER rule ensures that the bodies of the success clause and the operation clauses all have the output type D . The type of x in the value clause must match the input type C . The type of the parameter $p_\ell (A_\ell)$ and resumption $r_\ell (B_\ell \rightarrow D)$ in the operation clause H^ℓ is determined by the signature for ℓ . The return type of r_ℓ is D , as the body of the resumption will itself be handled by H . We write H^ℓ and H^{val} for projecting success and operation clauses.

$$\begin{aligned} H^\ell &:= \{ \ell p r \mapsto M \}, & \text{where } \{ \ell p r \mapsto M \} &\in H \\ H^{\text{val}} &:= \{ \text{val } x \mapsto M \}, & \text{where } \{ \text{val } x \mapsto M \} &\in H \end{aligned}$$

We extend the operational semantics to λ_h .

$$\begin{aligned} \text{handle } (\text{return } V) \text{ with } H &\rightsquigarrow N[V/x], \\ &\text{where } H^{\text{val}} = \{ \text{val } x \mapsto N \} \\ \text{handle } \mathcal{E}[\text{do } \ell V] \text{ with } H &\rightsquigarrow \\ N[V/p, \lambda y. \text{handle } \mathcal{E}[\text{return } y] \text{ with } H/r], & \\ &\text{where } H^\ell = \{ \ell p r \mapsto N \} \\ \mathcal{H}[M] &\rightsquigarrow \mathcal{H}[N], \quad \text{if } M \rightsquigarrow N \\ \text{Handler contexts } \mathcal{H} &::= [] \mid \text{handle } \mathcal{H} \text{ with } H \\ &\mid \text{let } x \leftarrow H \text{ in } N \end{aligned}$$

The first rule invokes the success clause. The second rule handles an operation via the corresponding operation clause. The third rule replaces the corresponding lifting rule for λ_b . Rather than augmenting evaluation contexts from λ_b , we introduce handler contexts. The separation between pure evaluation contexts \mathcal{E} and handler contexts \mathcal{H} guarantees

the second rule is deterministic, as otherwise it could pick an arbitrary handler in scope. With this separation, the second rule always picks the innermost handler.

We now characterise normal forms and state the standard type soundness property of λ_h .

Definition 3.1 (Computation normal forms). We say that a computation term N is normal with respect to $\ell \in \Sigma$, if N is either of the form **return** V , or $\mathcal{E}[\text{do } \ell W]$.

Theorem 3.2 (Type Soundness). *If $\vdash M : C$, then either there exists $\vdash N : C$ such that $M \rightsquigarrow^+ N$ and N is normal, or M diverges.*

4 Abstract Machine Semantics

Thus far we have introduced the base calculus λ_b and its extension with effect handlers λ_h . For each calculus we have given a *small-step operational semantics* which uses a substitution model for evaluation. Whilst this model is semantically pleasing, it falls short of providing a realistic account of practical computation as substitution is an expensive operation. We now develop a more practical model of computation based on an *abstract machine semantics*.

4.1 Base Machine

We choose a CEK-style abstract machine semantics [14] for λ_b based on that of Hillerström and Lindley [18]. The CEK machine operates on configurations which are triples of the form $\langle M \mid \gamma \mid \sigma \rangle$. The first component contains the computation currently being evaluated. The second component

496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550

contains the environment γ which binds free variables. The third component contains the continuation which instructs the machine how to proceed once evaluation of the current computation is complete. The syntax of abstract machine states is as follows.

Configurations	$C \in \text{Conf} ::= \langle M \mid \gamma \mid \sigma \rangle$
Environments	$\gamma \in \text{Env} ::= \emptyset \mid \gamma[x \mapsto v]$
Machine values	$v, w \in \text{Mval} ::= x \mid n \mid c \mid \langle \rangle \mid \langle v, w \rangle$ $\mid (\gamma, \lambda x^A. M) \mid (\gamma, \mathbf{rec} f x^A. M)$ $\mid (\mathbf{inl} v)^B \mid (\mathbf{inr} w)^A$
Continuations	$\sigma \in \text{PureCont} ::= [] \mid (\gamma, x, N) :: \sigma$

Values consist of function closures, constants, pairs, and left or right tagged values. A continuation is a stack of continuation frames. A continuation frame (γ, x, N) closes a let-binding $\mathbf{let} x \leftarrow [] \mathbf{in} N$ over environment γ . We write $[]$ for an empty continuation and $\phi :: \sigma$ for the result of pushing the frame ϕ onto σ . We use pattern matching to deconstruct continuations.

The abstract machine semantics is given in Figure 3. The transition function is given by the \longrightarrow relation, which also depends on an interpretation function $\llbracket - \rrbracket$ for value terms and machine values. The machine is initialised by placing a term in a configuration alongside the empty environment (\emptyset) and identity continuation ($[]$). The rules (M-APP), (M-REC), (M-CONST), (M-SPLIT), (M-CASEL), and (M-CASER) eliminate values. The (M-LET) rule extends the current continuation with let bindings. The (M-RETCONT) rule binds a returned value if there is a pure continuation in the current continuation frame. Given an input of a well-typed closed computation term $\vdash M : A$, the machine will either diverge or return a value of type A . A final state is given by a configuration of the form $\langle \mathbf{return} V \mid \gamma \mid [] \rangle$ in which case the final return value is given by the denotation $\llbracket V \rrbracket \gamma$ of V under environment γ . We now make the correspondence between operational semantics and abstract machine more precise.

Correctness The abstract machine faithfully simulates the operational semantics; most transitions correspond directly to β -reductions, but the M-LET-rule performs an administrative step to bring the computation M into evaluation position. We define an extension of the transition function \longrightarrow to capture administrative steps.

Definition 4.1 (Auxiliary reduction relations). We write \longrightarrow_a for administrative steps, \longrightarrow_β for β -steps, and \Longrightarrow for a sequence of steps of the form $\longrightarrow_a^* \longrightarrow_\beta$.

Theorem 4.2. *There is a one-to-one mapping between reduction relation \rightsquigarrow and transition function \Longrightarrow .*

The proof is by induction on $M \rightsquigarrow N$, relying on an inverse map $(-)$ from configurations to terms [18].

4.2 Handler Machine

We now enrich the λ_b machine to a λ_h machine. To support handlers we extend the syntax as follows.

Configurations	$C \in \text{Conf} ::= \langle M \mid \gamma \mid \kappa \rangle$
Continuations	$\kappa \in \text{Cont} ::= [] \mid (\sigma, \chi) :: \kappa$
Handler closures	$\chi \in \text{HClo} ::= (\gamma, H)$
Machine values	$v, w \in \text{Mval} ::= \dots \mid \chi$

The notion of configurations changes slightly as the continuation component is now occupied by a generalised continuation $\kappa \in \text{Cont}$ [18], meaning a machine continuation is now a list of pairs containing a pure continuation (as in the base machine) and a handler closure (χ). A handler closure consists of an environment and a handler definition, where the former binds the free variables that occur in the latter. The identity continuation is an empty pure continuation paired with the identity handler closure, i.e. $\kappa_0 := [([], (\emptyset, \{\mathbf{val} x \mapsto x\}))]$. The machine values are augmented to contain handler closures as an operation invocation causes the topmost handler closure of the machine continuation to be reified (and bound to the resumption parameter in the operation clause).

The handler machine adds transition rules for handlers, and modifies (M-LET) and (M-RETCONT) from the base machine to account for the richer continuation structure. Figure 4 depicts the new and modified rules. The rule (M-HANDLE) pushes a the handler closure along with an empty pure continuation onto the continuation stack. The (M-LET) and (M-RETCONT) are dual rules, as the former grows the pure continuation of the topmost continuation frame, and the latter shrinks the pure continuation. If the pure continuation is empty, then the (M-RETHANDLER) rule applies, which transfers control to the success clause of the current handler. If an operation is invoked, then the (M-HANDLE-OP) rule transfers control to the corresponding operation clause on the topmost handler and during the process it reifies the handler closure. Finally, the (M-RESUME) rule applies a reified handler closure, by pushing it onto the continuation stack. The handler machine has two possible final states: either it yields a value or it gets stuck on an unhandled operation.

Correctness Theorem 4.2 can mostly be repurposed for the handler machine as we need only recheck the cases for (M-LET) and (M-RETCONT) and check the cases for handlers.

4.3 Realisability and Asymptotic Complexity

The machine structures are readily realisable using standard persistent functional data structures. The pure and generalised continuations can be implemented using lists, which makes their augmentation operation $(_ :: _)$ have time complexity $O(1)$. This also holds true for pure continuations on the handler machine as augmenting the current pure continuation only requires reaching under the topmost handler closure. Environments, γ , can be realised using a map, making the complexity of extension and lookup be $O(\log |\gamma|)$ [32].

Transition function

661	M-APP	$\langle V W \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma' [x \mapsto \llbracket W \rrbracket \gamma] \mid \sigma \rangle,$	if $\llbracket V \rrbracket \gamma = (y', \lambda x^A. M)$	716		
662	M-REC	$\langle V W \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma' [f \mapsto (y', \mathbf{rec} f x.M), x \mapsto \llbracket W \rrbracket \gamma] \mid \sigma \rangle,$	if $\llbracket V \rrbracket \gamma = (y', \mathbf{rec} f x^A.M)$	717		
663	M-CONST	$\langle V W \mid \gamma \mid \sigma \rangle \longrightarrow \langle \mathbf{return} (\ulcorner c \urcorner (\llbracket V \rrbracket \gamma)) \mid \gamma \mid \sigma \rangle,$	if $\llbracket V \rrbracket \gamma = c$	718		
664	M-SPLIT	$\langle \mathbf{let} (x, y) = V \mathbf{in} N \mid \gamma \mid \sigma \rangle \longrightarrow \langle N \mid \gamma [x \mapsto v, y \mapsto w] \mid \sigma \rangle,$	if $\llbracket V \rrbracket \gamma = \langle v; w \rangle$	719		
665	M-CASEL	$\langle \mathbf{case} V \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \} \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma [x \mapsto v] \mid \sigma \rangle,$	if $\llbracket V \rrbracket \gamma = \mathbf{inl} v$	720		
666	M-CASER	$\langle \mathbf{case} V \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \} \mid \gamma \mid \sigma \rangle \longrightarrow \langle N \mid \gamma [y \mapsto v] \mid \sigma \rangle,$	if $\llbracket V \rrbracket \gamma = \mathbf{inr} v$	721		
667	M-LET	$\langle \mathbf{let} x \leftarrow M \mathbf{in} N \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma \mid (\gamma, x, N) :: \sigma \rangle$		722		
668	M-RETCONT	$\langle \mathbf{return} V \mid \gamma \mid (\gamma', x, N) :: \sigma \rangle \longrightarrow \langle N \mid \gamma' [x \mapsto \llbracket V \rrbracket \gamma] \mid \sigma \rangle$		723		
669	Value interpretation					
670				724		
671	$\llbracket x \rrbracket \gamma = \gamma(x)$	$\llbracket n \rrbracket \gamma = n$	$\llbracket \lambda x^A. M \rrbracket \gamma = (\gamma, \lambda x^A. M)$	$\llbracket \langle V; W \rangle \rrbracket \gamma = \langle \llbracket V \rrbracket \gamma; \llbracket W \rrbracket \gamma \rangle$	$\llbracket (\mathbf{inl} V)^B \rrbracket \gamma = (\mathbf{inl} \llbracket V \rrbracket \gamma)^B$	725
672	$\llbracket \langle \rangle \rrbracket \gamma = \langle \rangle$	$\llbracket c \rrbracket \gamma = c$	$\llbracket \mathbf{rec} f x^A. M \rrbracket \gamma = (\gamma, \mathbf{rec} f x^A. M)$	$\llbracket (\mathbf{inr} V)^A \rrbracket \gamma = (\mathbf{inr} \llbracket V \rrbracket \gamma)^A$		726
673					727	
674					728	
675					729	
676					730	
677					731	
678					732	
679					733	
680					734	
681					735	
682					736	
683					737	
684					738	
685					739	
686					740	
687					741	
688					742	
689					743	
690					744	
691					745	
692					746	
693					747	
694					748	
695					749	
696					750	
697					751	
698					752	
699					753	
700					754	
701					755	
702					756	
703					757	
704					758	
705					759	
706					760	
707					761	
708					762	
709					763	
710					764	
711					765	
712					766	
713					767	
714					768	
715					769	

Figure 3. Abstract Machine Semantics for λ_b **Transition function**

677	M-RESUME	$\langle V W \mid \gamma \mid \kappa \rangle \longrightarrow \langle \mathbf{return} W \mid \gamma \mid (\sigma, \chi) :: \kappa \rangle,$	if $\llbracket V \rrbracket \gamma = (\sigma, \chi)^A$	732
678	M-LET	$\langle \mathbf{let} x \leftarrow M \mathbf{in} N \mid \gamma \mid (\sigma, \chi) :: \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ((\gamma, x, N) :: \sigma, \chi) :: \kappa \rangle$		733
679	M-RETCONT	$\langle \mathbf{return} V \mid \gamma \mid ((\gamma', x, N) :: \sigma, \chi) :: \kappa \rangle \longrightarrow \langle N \mid \gamma' [x \mapsto \llbracket V \rrbracket \gamma] \mid (\sigma, \chi) :: \kappa \rangle$		734
680	M-HANDLE	$\langle \mathbf{handle} M \mathbf{with} H \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma \mid (\llbracket \cdot \rrbracket, (\gamma, H)) :: \kappa \rangle$		735
681	M-RETHANDLER	$\langle \mathbf{return} V \mid \gamma \mid (\llbracket \cdot \rrbracket, (\gamma', H)) :: \kappa \rangle \longrightarrow \langle M \mid \gamma' [x \mapsto \llbracket V \rrbracket \gamma] \mid \kappa \rangle,$	if $H^{\text{val}} = \{ \mathbf{val} x \mapsto M \}$	736
682	M-HANDLE-OP	$\langle \mathbf{do} \ell V \mid \gamma \mid (\sigma, (\gamma', H)) :: \kappa \rangle \longrightarrow \langle M \mid \gamma' [p \mapsto \llbracket V \rrbracket \gamma, r \mapsto (\sigma, (\gamma', H))] \mid \kappa \rangle,$	if $\ell : A \rightarrow B \in \Sigma$ and $H^\ell = \{ \ell p r \mapsto M \}$	737
683				738
684				739
685				740
686				741
687				742
688				743
689				744
690				745
691				746
692				747
693				748
694				749
695				750
696				751
697				752
698				753
699				754
700				755
701				756
702				757
703				758
704				759
705				760
706				761
707				762
708				763
709				764
710				765
711				766
712				767
713				768
714				769
715				770

Figure 4. Abstract Machine Semantics for λ_h

The worst-case time complexity of the machine transition relation \longrightarrow is exhibited by rules which involve operations on the environment, since any other operation is constant time, hence the worst-time complexity of a transition is $O(\log |\gamma|)$. The value interpretation function $\llbracket - \rrbracket \gamma$ is defined structurally on values. Its worst-time complexity is exhibited by a nesting of pairs of variables $\llbracket \langle x_1, \dots, x_n \rangle \rrbracket \gamma$ which has complexity $O(n \log |\gamma|)$.

Continuation copying On the handler machine the top-most continuation frame can be copied in constant time due to the persistent runtime and the layout of machine continuation. An alternative design would be to make the runtime non-persistent in which case copying a continuation frame $((\sigma, (\gamma, _)) :: _)$ would be a $O(|\sigma| + |\gamma|)$ time operation.

5 Efficient Generic Search

We now come to the crux of the paper. In this section we prove that λ_h accommodates some programmable operations with an asymptotic runtime bound that cannot be achieved in λ_b . Since addition of effect handlers is the only difference between the two languages, we obtain as a corollary that a PCF-like programming language with effect handlers exhibits fundamentally more efficient programs than a pure PCF programming language. To obtain this result it suffices to find *one* efficient program in λ_h and show that *no* equivalent program in λ_b can achieve the same asymptotic complexity: we take *generic search*.

Generic search is a modular search procedure that finds solutions to a given search problem P . Generic search is agnostic to the specific instantiation of P , and as a result is applicable across a wide spectrum of domains. A variety of problems can be cast as instances of generic search; classic examples include solving Sudoku puzzles and n -Queens, whilst more esoteric examples include problems from game theory, graph theory, and exact real number integration [10, 39].

To simplify the presentation, we compute the number of solutions (generic count), rather than materialising all solutions (generic search). With little extra effort one can tweak the development to compute exact solutions.

Informally, a generic count program takes as input a predicate and returns the number of times the predicate yields true. A predicate returns a boolean value which signifies whether its input satisfies the predicate. As input a predicate takes a bit vector of length $n > 0$, which we represent as a first-order function $\text{Nat} \rightarrow \text{Bool}$. Ultimately we ask for implementations of a program, count, whose type is

$$\text{count}_n : ((\text{Nat}_n \rightarrow \text{Bool}) \rightarrow \text{Bool}) \rightarrow \text{Nat}$$

where Nat_n admits elements of the set $\mathbb{N}_n := \{0, \dots, n-1\}$. We often omit n indexes when clear from context; in particular they do not appear explicitly in the types of our programs as our formalism does not support dependent types.

Before giving the necessary formal machinery to state and prove the result, we first introduce the concepts informally.

5.1 Predicates and Points

Higher-order functions are the key to our modular formulation of generic search. We define a predicate of size n as a higher-order function which acts on points

$$\text{Predicate}_n := \text{Point}_n \rightarrow \text{Bool}$$

where n is a natural number and a point is a first-order function taking bounded natural numbers to boolean values:

$$\text{Point}_n := \text{Nat}_n \rightarrow \text{Bool}$$

Intuitively, a point implements a vector of boolean values where the natural number argument serves as an index into the vector. A point need not be a total function; indeed points we concern ourselves with are typically partial.

Examples Let us consider some simple examples of predicates and points. As a first example consider the constant point, $p_{\text{true}} := \lambda_.\text{true}$. A slightly more interesting point is

$$p_2 := \lambda i. \text{if } i = 0 \text{ then true else if } i = 1 \text{ then false else } \perp i$$

where $\perp := \text{rec } f \ i.f \ i$ is the always-diverging point.

Now let us move onto some example predicates. We can give a whole family of constant true predicates. For example tt_0 returns true irrespective of its point.

$$\text{tt}_0 := \lambda p. \text{true}$$

Similarly we can define a variation, tt_2 , which inspects two components of its point, but nevertheless returns true.

$$\text{tt}_2 := \lambda p. p\ 1; p\ 0; \text{true}$$

This predicate is slightly more interesting than tt_0 as it is defined only for points defined on Nat_n for $n \geq 2$. A predicate may inspect the same component of its point more than once

$$\text{red}_1 := \lambda p. p\ 0; p\ 0$$

thus performing redundant work. Another class of predicates are divergent predicates such as

$$\text{div}_1 := \text{rec } \text{div } p. \text{if } p\ 0 \text{ then } \text{div } p \text{ else false}$$

which diverges whenever the 0-th index of the point yields true. Thus both $\text{div}_1\ p_{\text{true}}$ and $\text{div}_1\ p_2$ never terminate. Finally, let us consider a productive predicate which determines whether a point contains an odd number of true components.

$$\text{odd}_n := \lambda p. \text{fold } \otimes \ \text{false} \ (\text{map } p \ [0, \dots, n-1])$$

where fold and map are the standard combinators on lists, and \otimes is the exclusive-or function. This predicate is only well-defined for $n > 0$. Applying odd_2 to p_2 yields true, whereas applying it to p_{true} yields false.

Predicate Models In essence a predicate is a decision procedure, which participates in a ‘dialogue’ with a supplied point $p : \text{Point}_n$. The predicate may *query* (i.e. invoke) the components of p , and p then *responds* (i.e. returns). Ultimately this dialogue may *answer* whether the point satisfies the predicate. We can model the behaviour of a predicate as an unrooted binary decision tree characterising the predicate’s interaction with p , where each interior node is labelled with a query $?i$ (for $i \in \mathbb{N}_n$) whose the left subtree corresponds




Figure 5. Example Decision Tree Models

to $p\ i$ being true and the right subtree to $p\ i$ being false, and each leaf is labelled with an answer $!\text{true}$ or $!\text{false}$ according to whether p satisfies the predicate. The trees are unrooted to account for the bit of computation that occurs in between the application of a predicate to p and the first query or answer.

Figure 5 depicts models of some of the example predicates given above. The model of tt_0 is simply an unrooted leaf (Figure 5a). The model of div_1 is an infinite left-branching tree (Figure 5b). The model of odd_2 is a complete binary tree (Figure 5c). A further example is the unconditionally divergent predicate $\text{div} := \text{rec } f \ p.f \ p$ whose model is empty.

Restrictions To obtain a meaningful complexity result, we must constrain the predicates of interest. At one extreme, counting the size of a divergent predicate like div_0 is meaningless. At the other extreme, a constant predicate like tt_0 exhibits no interesting computational characteristics; other constant predicates like tt_2 inspect their provided point. Predicates like red_1 perform redundant work. Such redundancy can be eliminated via insertion of a local let binding.

Thus we restrict attention to predicates that for $n > 0$

1. terminate when applied to any point p ; and
2. inspect each bit $0 < i < n$ of p exactly once.

Of the examples so far, the ones satisfying the conditions are tt_2 and odd_n . Predicates satisfying 1 and 2 are exactly those whose models form complete binary trees (as in Figure 5c), which we call *n-standard*. We provide a rigorous definition of *n-standard* predicates in Section 5.3. To satisfy 1, we also require that points terminate on their defined domain Nat_n . We call a point that is defined on $0 < i < n$ an *n-point*.

5.2 Effectful Generic Counting

Having introduced predicates and points informally, we move onto presenting our effectful implementation of count. Our implementation is a variation of the example handler for non-deterministic computation that we gave in Section 2. The main idea is to implement points as non-deterministic computations using the Branch operation such that the handler may respond to every query twice by invoking the provided resumption with true and subsequently false. The key insight is that the resumption restarts computation at the invocation site of Branch, which means that prior computation need not be repeated. In other words, the resumption

ensures that common bits of computations prior to any query are shared between both branches.

We fix the effect signature $\Sigma := \{\text{Branch} : 1 \rightarrow \text{Bool}\}$. The algorithm is then expressed as follows.

```

885   effcount : ((Nat → Bool) → Bool) → Nat
886   effcount P :=
887     handle P (λ_.do Branch ⟨⟩) with
888       val b      ↦ if b then return 1 else return 0
889       Branch ⟨ r ↦ let xtrue ← r true in
890                   let xfalse ← r false in xtrue + xfalse

```

The handler applies predicate P to a single point defined using `Branch`. The boolean return value is interpreted as a single solution, whilst `Branch` is interpreted by alternately supplying true and false to the resumption and summing the results. A curious detail about `effcount` is that it works for all n -standard predicates without having to know the exact value of n . This is because the point $(\lambda_.\text{do Branch } \langle \rangle)$ represents the superposition of all possible points. The sharing enabled by the use of the resumption is exactly the ‘magic’ we need to make it possible to implement generic counting more efficiently in λ_h than in λ_b .

5.3 Predicates, Points, and their Models, Formally

We now formalise the notions of n -standard predicates, points, and their models. We formalise the concepts using the operational semantics and abstract machine for the base language λ_b . The reason being that it makes little sense to compare the runtime complexity of predicates which makes use effectful operations as they cannot be run on the base machine.

We begin by formalising the decision tree model of predicates. We first introduce the label set, `Lab`, consisting of queries and answers.

Notation. We write $bs \sqsubset bs'$ to mean that list bs is a prefix of list bs' .

Definition 5.1 (label set). The label set `Lab` consists of queries parameterised by a natural number and answers parameterised by a boolean.

$$\text{Lab} := \{?n \mid n \in \mathbb{N}\} \cup \{!true, !false\}$$

We model a decision tree as a partial function from lists of booleans to labels; each boolean list specifies a cursor into the tree as a path from the root of the tree.

Definition 5.2 (untimed) decision tree). A decision tree is a partial function $t : \mathbb{B}^* \rightarrow \text{Lab}$ from lists of booleans to node labels with the following properties:

- The domain of t , $\text{dom}(t)$, is prefix closed.
- If $t(bs) = !b$ then $t(bs')$ is undefined for all $bs' \sqsupset bs$. In other words answer nodes are always leaves.

Timed decision trees are decorated with timing data that records the number of machine steps.

Definition 5.3 (timed decision tree). A timed decision tree is a partial function $t : \mathbb{B}^* \rightarrow \text{Lab} \times \text{Nat}$ such that its first

projection $bs \mapsto t(bs).1$ is a decision tree. We write $\text{labs}(t)$ for the first projection ($bs \mapsto t(bs).1$) and $\text{steps}(t)$ for the second projection ($bs \mapsto t(bs).2$) of a timed decision tree.

We now relate predicates to decision trees by way of an interpretation of configurations as decision trees.

Notation. We write $a \simeq b$ for Kleene equality: either both a and b are undefined or both are defined and $a = b$.

Definition 5.4. The timed decision tree of a configuration is defined by the following equations

$$\begin{aligned} \mathcal{T}(\langle \text{return true} \mid \gamma \mid [] \rangle) [] &= (!\text{true}, 0) \\ \mathcal{T}(\langle \text{return false} \mid \gamma \mid [] \rangle) [] &= (!\text{false}, 0) \\ \mathcal{T}(\langle p V \mid \gamma \mid \sigma \rangle) [] &= (?[V]\gamma, 0) \end{aligned}$$

$$\begin{aligned} \mathcal{T}(\langle p V \mid \gamma \mid \sigma \rangle) (b :: bs) &\simeq \mathcal{T}(\langle \text{return } b \mid \gamma \mid \sigma \rangle) bs \\ \mathcal{T}(\langle M \mid \gamma \mid \sigma \rangle) bs &\simeq \mathcal{I}(\mathcal{T}(\langle M' \mid \gamma' \mid \sigma' \rangle) bs), \\ &\text{if } \langle M \mid \gamma \mid \sigma \rangle \longrightarrow \langle M' \mid \gamma' \mid \sigma' \rangle \end{aligned}$$

where $\mathcal{I}(\ell, s) = (\ell, s+1)$ and p is a distinguished free variable such that in all of the above equations $\gamma(p) = \gamma'(p) = p$. The decision tree of a computation term is obtained by placing it in the initial configuration, i.e. $\mathcal{T}(M) := \mathcal{T}(\langle M, \emptyset[p \mapsto p], \kappa_0 \rangle)$. The decision tree of a predicate P is $\mathcal{T}(Pp)$. Since p is a parametric variable, we shall omit p and simply write $\mathcal{T}(P)$ to mean $\mathcal{T}(Pp)$.

We can define a construction procedure, \mathcal{U} , for untimed decision trees using \mathcal{T} as follows: $\mathcal{U}(P) := bs \mapsto \mathcal{T}(P)(bs).1$.

Definition 5.5 (n -standard trees and n -standard predicates). For any $n > 0$ a decision tree t is said to be n -standard if

- The domain of t consists of all the lists whose length is at most n , i.e., $\text{dom}(t) = \{bs : \mathbb{B}^* \mid |bs| \leq n\}$.
- Every leaf node in t is an answer node, i.e., for all $bs \in \text{dom}(t)$ if $|bs| = n$ then $t(bs) = !b$, for some $b \in \mathbb{B}$.
- There are no repeated queries along any path in t : for all $bs, bs' \in \text{dom}(t), j \in \mathbb{N}$, if $bs \sqsubseteq bs'$ and $t(bs) = ?j$ then $t(bs') = ?j$.

A timed decision tree t is n -standard if its underlying untimed decision tree ($bs \mapsto t(bs).1$) is. A predicate P is said to be n -standard if its decision tree $\mathcal{T}(P)$ is an n -standard tree.

As alluded to in Section 5.1 n -standard decision tree models are exactly those that form a complete binary tree such that each path contains no repeated queries. The third condition in the definition requires only that there are no repeated queries along any path in the model; it does not impose a particular ordering on those queries.

We now move onto formalising points. Our model of points is only used for extensional reasoning about programs in the λ_b -language as we can reason intensionally about the single point used by `effcount` in the λ_h -language. As remarked in Section 5.1, points may in general be partial, however, the points that we shall consider all have the property, that they terminate whenever applied to an element of their defined domain (Nat_n for some $n > 0$).

Notation. We write $(-)_n : \mathbb{N} \rightarrow \text{Nat}$ for the injection of natural numbers into value terms and $\mathbb{N}[-] : \text{Nat} \rightarrow \mathbb{N}$ for its inverse. Similarly, we write $\mathbb{B}[-] : \text{Bool} \rightarrow \mathbb{B}$ for the denotation function for boolean value terms.

Definition 5.6 (n -point). For any $n > 0$ a closed value $p : \text{Point}_n$ is said to be an n -point if

$$\forall i \in \mathbb{N}_n. p(i) \rightsquigarrow^* \text{return } W.$$

Semantically, we can think of any n -point as a total first-order function of type $\mathbb{N}_n \rightarrow \mathbb{B}$. In fact, we shall take this function type to be the model of n -points. Since an n -point terminates on its defined domain, we can easily compute a model of it using the operational semantics. Definition 5.7 provides a procedure for computing the model of any n -point.

Definition 5.7. The denotation of an n -point p is the realisation of its operational behaviour

$$\begin{aligned} \mathbb{P}[-] &: (\text{Nat}_n \rightarrow \text{Bool}) \rightarrow (\mathbb{N}_n \rightarrow \mathbb{B}) \\ \mathbb{P}[p] &:= j \in \mathbb{N}_n \mapsto \mathbb{B}[p(j)] \end{aligned}$$

Definition 5.8. Any two n -points p_0 and p_1 are *distinct* if their denotations differ, i.e. $\exists j \in \mathbb{N}_n. \mathbb{P}[p_0] j \neq \mathbb{P}[p_1] j$.

5.4 Specification of Generic Counting

We now formally define generic counting.

Definition 5.9. A counting function is a partial function of type $\mathbb{B}^* \rightarrow \mathbb{N}$.

As with the decision tree functions, the list argument to a counting function serves as a cursor into the model of the predicate. However, in this case, the function computes the sum of the true answers in the subtree pointed to by the cursor. Thus in order to compute the sum of all true answers we apply the counting function to the empty list. The following definition provides a procedure for constructing a counting function for any predicate.

Definition 5.10. The counting function for a configuration is defined by the following equations.

$$\begin{aligned} C(\langle \text{return true} \mid \gamma \mid [] \rangle) &= 1 \\ C(\langle \text{return false} \mid \gamma \mid [] \rangle) &= 0 \\ C(\langle p V \mid \gamma \mid \sigma \rangle) &= C(\langle \text{return true} \mid \gamma \mid \sigma \rangle) + C(\langle \text{return false} \mid \gamma \mid \sigma \rangle) \\ C(\langle p V \mid \gamma \mid \sigma \rangle)(b :: bs) &\simeq C(\langle \text{return } b \mid \gamma \mid \sigma \rangle) bs \\ C(\langle M \mid \gamma \mid \sigma \rangle) bs &\simeq C(\langle M' \mid \gamma' \mid \sigma' \rangle) bs, \\ &\text{if } \langle M \mid \gamma \mid \sigma \rangle \longrightarrow \langle M' \mid \gamma' \mid \sigma' \rangle \end{aligned}$$

where p is a distinguished free variable such that in all of the above equations $\gamma(p) = \gamma'(p) = p$. As with \mathcal{T} , we write $C(P)$ for $C(P p)$.

Definition 5.11 (generic count program). A program $C : ((\text{Nat} \rightarrow \text{Bool}) \rightarrow \text{Bool}) \rightarrow \text{Nat}$ is said to be an n -count program if for every n -standard predicate P

$$C P \rightsquigarrow^+ \text{return } V, \text{ such that } \llbracket V \rrbracket = C(P)(\llbracket [] \rrbracket)$$

5.5 Complexity of Effectful Generic Counting

In this section we formulate correctness and asymptotic bounds for running the effectful generic counting program `effcount` on a predicate P . Full proofs are in Appendix B.

A key feature of the proof is that we must alternate between intensional and extensional modes of reasoning. As `effcount` is a fixed program, we can reason intensionally about its behaviour and thereby directly observe machine transitions. But we must also consider the transitions of P . Since the code for P is unknown we cannot employ the same reasoning technique. Instead, we reason extensionally by making use of the fact that the timed decision tree model of P contains the exact number of transitions that P performs in each branch of computation.

Theorem 5.12. For all $n > 0$ and any n -standard predicate P it holds that

1. The program `effcount` is a generic counting program: $\text{effcount } P \rightsquigarrow^+ \text{return } V$ such that $\mathbb{N}[\llbracket V \rrbracket] = C(P)(\llbracket [] \rrbracket) \leq 2^n$
2. The runtime complexity of `effcount` P is given by:

$$\sum_{bs \in \mathbb{B}^*}^{|bs| \leq n} \text{steps}(\mathcal{T}(P))(bs) + O(2^n)$$

5.6 Pure Generic Counting

We have shown that there exists an implementation of `count` in λ_h with a particular runtime bound. We now show that no implementation of `count` in λ_b can match this bound. To do so we exhibit two properties of the decision model:

1. there are no shortcuts, i.e. every leaf must be visited;
2. there can be no sharing of work amongst branches.

Together these two properties imply that every count program has least time complexity $\Omega(n2^n)$, because it must construct 2^n points, one for each leaf in the model, and apply the predicate once to each point. Due to the lack of sharing, each application of the predicate performs some redundant work as the path to two neighbouring leaves share n edges in the model. We formalise the first property in Section 5.7 and the second in Section 5.8. First, we give an example of a pure generic count program and discuss better alternatives.

The following is a direct implementation of `count` in λ_b .

```

purecountn : ((Natn → Bool) → Bool) → Nat
purecountn := λP. count' n ⊥
  where count' 0      p := if P p then 1 else 0
        count' (1 + n') p :=
          count' n' (λi. if i = n' then true else p i)
          + count' n' (λi. if i = n' then false else p i)
        ⊥ _           := rec f i f i

```

The implementation materialises 2^n points which are encoded using a standard functional linked list representation. The auxiliary function `count'` exhibits a recursion pattern reminiscent of the classic recursive definition of the Fibonacci function. The function is initially applied to the

divergent function \perp , which is partly used to seed the list generation, but also used to respond to queries $i \geq n$ such that they diverge. The base case of `count'` applies the predicate P to the generated point.

At this point the reader may wonder why we cannot simply use known continuation passing style (CPS) [19] or monadic [24] transforms of effect handlers, or implement an interpreter for effect handlers [18] in λ_b to achieve the sharing of computation. Such global implementation techniques are ruled out in our setting, because they would change the type of `count`. For example, as any predicate P is a higher-order function (supplied externally to `count`), we cannot even CPS transform `count` locally as the interface of P would be incompatible with the CPS interface. Many such transforms are possible in a first-order setting, but they are not an option for us due to the inherent higher-order nature of our setting.

5.7 No Shortcuts

We sketch the idea behind the proof of the fact that any n -count program in λ_b must construct 2^n points. Full details are in Appendix C.2. The proof makes use of the observation that the decision tree model encodes the canonical structure of any n -standard predicate. We can reify a (semantic) n -standard model as a (syntactic) n -standard predicate. This procedure provides a means for converting any n -standard predicate P into a canonical form (first compute the model, then reify, a la normalisation by evaluation [6]).

Lemma 5.13. *Suppose P is an n -standard predicate and C is an n -count program, then C applies P to at least 2^n distinct n -points.*

The lemma guarantees that any n -count program constructs an n -point corresponding to each leaf in the model of any given n -standard predicate.

5.8 No Sharing

We now show that distinct predicate applications cannot share computation. In order to do so, we introduce the notion of *threads*. Intuitively, a thread corresponds to a path in a decision tree model. Each thread of an n -standard model is composed of $n + 1$ *sections*, where each corresponds to an edge in the model. Thus we can identify the start and end of any section by looking for point queries and responses. Definition 5.15 makes use of an auxiliary reduction relation \rightarrow to define threads and sections. Intuitively, this reduction relation ensures that we cannot inadvertently “step over” an application of a point.

Definition 5.14 (\rightarrow). $\mathcal{E}[M] \rightarrow \mathcal{E}[N]$ iff $\mathcal{E}[M] \rightsquigarrow \mathcal{E}[N]$ and M is not of the form $p V$ where p is a point.

Definition 5.15 (Sections and threads). A section is a pair of computations, where the first component marks the start of the section and the second marks the end. For an n -standard predicate P , a thread of P consists of $n + 1$ sections. Given

a denotation, f , of a concrete n -point, and taking p to be a distinguished free variable, a single thread for P can be computed as follows

$$\begin{aligned} \text{Th} &: \text{Comp} \times (\mathbb{N}_n \rightarrow \mathbb{B}) \rightarrow [(\text{Comp}, \text{Comp})] \\ \text{Th}(P p, f) &:= (P p, \mathcal{E}[p V]) :: \text{Th}(\mathcal{E}[\mathbf{return} \langle b \rangle], f), \\ &\quad \text{where } \text{Sec}(P p) = \mathcal{E}[p V] \text{ and } b = f(\mathbb{N}[[V]]) \\ \text{Th}(\mathcal{E}[\mathbf{return} W], f) &:= \\ &\quad (\mathcal{E}[\mathbf{return} W], \mathcal{E}'[p V]) :: \text{Th}(\mathcal{E}[\mathbf{return} \langle b \rangle], f) \\ &\quad \text{where } \text{Sec}(\mathcal{E}[\mathbf{return} W]) = \mathcal{E}'[p V] \text{ and } b = f(\mathbb{N}[[V]]) \\ \text{Th}(\mathcal{E}[\mathbf{return} W], f) &:= (\mathcal{E}[\mathbf{return} W], \mathbf{return} V) :: [] \\ &\quad \text{where } \text{Sec}(\mathcal{E}[\mathbf{return} W]) = \mathbf{return} V \end{aligned}$$

The auxiliary procedure `Sec` computes the end of a section from the start.

$$\text{Sec}(\mathcal{E}[M]) := \begin{cases} \mathcal{E}'[p V] & \text{if } \mathcal{E}[M] \rightarrow^+ \mathcal{E}'[p V] \\ \mathbf{return} V & \text{if } \mathcal{E}[M] \rightarrow^* \mathbf{return} V \end{cases}$$

Now we show that every predicate application gives rise to a corresponding thread via $\text{Th}(-, -)$.

Lemma 5.16. *Suppose P is an n -standard predicate, p is an n -point, and $f = \mathbb{P}[[p]]$, then*

$$\begin{aligned} P p \rightarrow^+ \mathcal{E}_1[p V_1] \rightsquigarrow^+ \mathcal{E}_1[\mathbf{return} \langle f(\mathbb{N}[[V_1]]) \rangle] \rightarrow^+ \dots \\ \rightarrow^+ \mathcal{E}_n[p V_n] \rightsquigarrow^+ \mathcal{E}_n[\mathbf{return} \langle f(\mathbb{N}[[V_n]]) \rangle] \rightarrow^* \mathbf{return} W \end{aligned}$$

if and only if

$$\text{Th}(P p, f) = \left[\begin{array}{l} (P p, \mathcal{E}_1[p V_1]), \\ (\mathcal{E}_1[\mathbf{return} \langle f(\mathbb{N}[[V_1]]) \rangle], \mathcal{E}_2[p V_2]), \\ \vdots \\ (\mathcal{E}_n[p V_n], \mathcal{E}_n[\mathbf{return} \langle f(\mathbb{N}[[V_n]]) \rangle]), \\ (\mathcal{E}_n[\mathbf{return} \langle f(\mathbb{N}[[V_n]]) \rangle], \mathbf{return} W) \end{array} \right]$$

The lemma tells us that every predicate application has an associated thread and vice versa. By Lemma 5.13 we know that any n -count program must construct at least 2^n distinct threads. To establish the desired result, we need some way of characterising disjointness of threads.

Definition 5.17. Let C denote an n -count program and P an n -standard predicate. Any two threads T_0 and T_1 arising from distinct applications of P in C are said to be disjoint if every section computation of T_0 is distinct from every section computation of T_1 , where the section computations of a thread comprise the set of all start and end computations of the sections in that thread.

Now we may conclude that no two distinct predicate applications can share computation, or in other words every section of their associated threads must be evaluated.

Lemma 5.18. *Suppose P is an n -standard predicate and C is an n -count program, and let p_0 and p_1 be distinct n -points, then the predicate applications $P p_0$ and $P p_1$ within C have disjoint threads.*

Parameter	Queens						Integration								
	First solution			All solutions			Id	Squaring			Logistic				
	20	24	28	8	10	12		20	14	17	20	1	2	3	4
Naïve	∞	∞	∞	274.18	∞	∞	17.17	50.61	65.8	80.58	∞	∞	∞	∞	∞
Berger	9.29	12.69	∞	2.11	2.81	3.41	5.59	23.30	25.65	27.50	26.10	33.27	34.02	32.76	31.00
Pruned	2.03	2.37	2.66	1.29	1.42	1.52	2.27	4.39	5.00	5.08	4.80	6.25	7.18	8.09	8.80
Bespoke	0.13	0.12	0.12	0.15	0.05	0.04									

Table 1. Speedup of the Effectful Implementation

5.9 Complexity of Pure Generic Counting

Now we can plug together the formal machinery developed in the previous sections to state and prove the complexity result for pure generic counting programs.

Theorem 5.19. *For all $n > 0$ and every n -count program $\text{count} \in \lambda_b$, and n -standard predicate the runtime of $\text{count } P$ is at least*

$$\sum_{\substack{|bs| \leq n \\ bs \in \mathbb{B}^*}} 2^{n-|bs|} \text{steps}(t)(bs) + \Omega(n2^n)$$

where $t = \mathcal{T}(P)$.

Proof. By composing Lemmas 5.13 and 5.18 we obtain the result in terms of the operational semantics. To lift this result to the abstract machine semantics we apply Theorem 4.2. \square

6 Robustness

Our complexity result is robust as it remains true in more general settings. We outline here how it generalises beyond n -standard predicates and to richer base languages.

Beyond n -standard Predicates The n -standard restriction ensures that the models of predicates are complete binary trees. It is possible to relax the restriction at the expense of a more complicated analysis. The restriction serves to make the result as crisp as possible. Nevertheless, we may augment the effcount program with internal state to keep track of the queries raised by predicates and to remember answers of previous queries. State is definable in λ_h using a standard technique known as parameter-passing [38].

Mutable State Mutable state is a staple ingredient of many practical programming languages. To support mutable state, the base language is endowed with a heap (giving rise to a CESK machine [12], where S means *store*). Modulo heap bookkeeping, the analysis is the same.

7 Experiments

The theoretical efficiency gap between realisations of λ_b and λ_h manifests in practice. We have observed it empirically on instantiations of n -Queens and exact real number integration, which can be cast as generic search. Table 1 shows the speedup of using an effectful implementation of generic

search over various pure implementations. We discuss the benchmarks and results in further detail below.

Methodology We evaluated an effectful implementation of generic search against three “pure” implementations which are realisable in λ_b extended with mutable state:

- Naïve: a variation of the purecount program;
- Pruned: a generic search procedure with space pruning based on Longley’s technique [28] (uses local state);
- Berger: a lazy pure functional generic search procedure based on Berger’s algorithm [5].

Each benchmark was run 11 times. The reported figure is the median runtime ratio between the particular implementation and the baseline effectful implementation. Benchmarks that failed to terminate within a threshold (1 minute for single solution, 8 minutes for enumerations), are reported as ∞ . The experiments were conducted in SML/NJ v110.78 with factory settings on an Intel Xeon CPU E5-1620 v2 @ 3.70GHz powered workstation running Ubuntu 16.04. The effectful implementation uses an encoding of delimited control akin to effect handlers based on top of SML/NJ’s call/cc.

Queens We phrase the n -Queens problem as a generic search problem. As a control we include a bespoke implementation hand-optimised for the problem. We perform two experiments: finding the first solution for $n \in \{20, 24, 28\}$ and enumerating all solutions for $n \in \{8, 10, 12\}$. The speedup over the naïve implementation is dramatic, but less so over the Berger procedure. The pruned procedure is more competitive, but still slower than the baseline. Unsurprisingly, the baseline is much slower than the bespoke implementation.

Exact Real Integration The integration benchmarks are adapted from Simpson [39]. We integrate three different functions with varying precision in the interval $[0, 1]$. For the identity function (Id) at precision 20 the speedup relative to Berger is $5.59\times$. For the squaring function the speedups are larger at higher precisions: at precision 14 the speedup is $4.39\times$ over the pruned integrator, whilst it is $5.08\times$ at precision 20. The speedups are more extreme against the naïve and Berger integrators. We also integrate the logistic map $x \mapsto 1 - 2x^2$ at a fixed precision of 15. We make the function harder to compute by iterating it up to 5 times.

Between the pruned and effectful integrator the speedup ratio increases as the function becomes harder to compute.

8 Conclusions and Future Work

We presented a PCF-inspired language λ_b and its extension with effect handlers λ_h . We proved that λ_h exhibits an asymptotically more efficient implementation of generic search than any possible implementation in λ_b . We observed its effect in practice on several benchmarks.

The result extends to other control operators by appeal to existing results on interdefinability of handlers and other control operators [17, 34]. The result no longer applies directly if we add an effect type system to λ_h , as the implementation of the counting program would require a change of type for predicates to reflect the ability to perform effectful operations. In future we plan to investigate how to account for effect type systems.

Acknowledgments

We would like to thank James McKinna for insightful discussions about this work. Daniel Hillerström was supported by EPSRC grant EP/L01503X/1 (EPSRC Centre for Doctoral Training in Pervasive Parallelism). Sam Lindley was supported by EPSRC grant EP/K034413/1 (From Data Types to Session Types—A Basis for Concurrency and Distribution).

References

[1] Andrej Bauer. 2011. How make the "impossible" functionals run even faster. Mathematics, Algorithms and Proofs, Leiden, the Netherlands. (2011). <http://math.andrej.com/2011/12/06/how-to-make-the-impossible-functionals-run-even-faster/>

[2] Andrej Bauer. 2018. What is algebraic about algebraic effects and handlers? *CoRR* abs/1807.05923 (2018).

[3] Andrej Bauer and Matija Pretnar. 2015. Programming with algebraic effects and handlers. *J. Log. Algebr. Meth. Program.* 84, 1 (2015), 108–123.

[4] Nick Benton and Andrew Kennedy. 2001. Exceptional Syntax Journal of Functional Programming. *J. Funct. Program.* 11, 4 (2001), 395–410.

[5] Ulrich Berger. 1990. *Totale Objekte und Mengen in der Bereichstheorie*. Ph.D. Dissertation. Ludwig Maximilians-Universität, Munich.

[6] Ulrich Berger, Matthias Eberl, and Helmut Schwichtenberg. 1998. Normalisation by Evaluation. In *Prospects for Hardware Foundations (Lecture Notes in Computer Science)*, Vol. 1546. Springer, 117–137.

[7] Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Abstracting algebraic effects. *PACMPL* 3, POPL (2019), 6:1–6:28.

[8] Richard Bird, Geraint Jones, and Oege de Moor. 1997. More haste less speed: lazy versus eager evaluation. *J. Funct. Progr.* 7, 5 (1997), 541–547.

[9] Robert Cartwright and Matthias Felleisen. 1992. Observable Sequentiality and Full Abstraction. In *POPL*. ACM Press, 328–342.

[10] Robbie Daniels. 2016. *Efficient Generic Searches and Programming Language Expressivity*. Master's thesis. School of Informatics, the University of Edinburgh, Scotland. http://homepages.inf.ed.ac.uk/jrl/Research/Robbie_Daniels_MSc_dissertation.pdf

[11] Stephen Dolan, Leo White, KC Sivaramakrishnan, Jeremy Yallop, and Anil Madhavapeddy. 2015. Effective Concurrency through Algebraic Effects. OCaml Workshop. (2015).

[12] Matthias Felleisen. 1987. *The Calculi of Lambda-nu-cs Conversion: A Syntactic Theory of Control and State in Imperative Higher-order Programming Languages*. Ph.D. Dissertation. Indianapolis, IN, USA. AA18727494.

[13] Matthias Felleisen. 1991. On the expressive power of programming languages. *Sci. Comput. Prog.* 17, 1–3 (1991), 35–75.

[14] Matthias Felleisen and Daniel P. Friedman. 1987. Control Operators, the SECD-machine, and the λ -Calculus. In *The Proceedings of the Conference on Formal Description of Programming Concepts III, Eberup, Denmark*. Elsevier, 193–217.

[15] Andrzej Filinski. 1994. Representing Monads. In *POPL*. ACM Press, 446–457.

[16] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. 1993. The Essence of Compiling with Continuations. In *PLDI*. ACM, 237–247.

[17] Yannick Forster, Ohad Kammar, Sam Lindley, and Matija Pretnar. 2017. On the expressive power of user-defined effects: effect handlers, monadic reflection, delimited control. *PACMPL* 1, ICFP (2017), 13:1–13:29.

[18] Daniel Hillerström and Sam Lindley. 2016. Liberating effects with rows and handlers. In *TyDe@ICFP*. ACM, 15–27.

[19] Daniel Hillerström, Sam Lindley, Robert Atkey, and K. C. Sivaramakrishnan. 2017. Continuation Passing Style for Effect Handlers. In *FSCD (LIPIcs)*, Vol. 84. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 18:1–18:19.

[20] Neil Jones. 2001. The expressive power of higher-order types, or, life without CONS. *J. Funct. Progr.* 11 (2001), 5–94.

[21] Ohad Kammar, Sam Lindley, and Nicolas Oury. 2013. Handlers in action. In *ICFP*. ACM, 145–158.

[22] Oleg Kiselyov, Amr Sabry, and Cameron Swords. 2013. Extensible effects: an alternative to monad transformers. In *Haskell*. ACM, 59–70.

[23] Donald Knuth. 1997. *The Art of Computer Programming, Volume 1: Fundamental Algorithms (third edition)*. Addison-Wesley.

[24] Daan Leijen. 2017. Type directed compilation of row-typed algebraic effects. In *POPL*. ACM, 486–499.

[25] Paul Blain Levy, John Power, and Hayo Thielecke. 2003. Modelling environments in call-by-value programming languages. *Inf. Comput.* 185, 2 (2003), 182–210.

[26] Sam Lindley. 2014. Algebraic effects and effect handlers for idioms and arrows. In *WGP@ICFP*. ACM, 47–58.

[27] Sam Lindley, Conor McBride, and Craig McLaughlin. 2017. Do be do be do. In *POPL*. ACM, 500–514.

[28] John Longley. 1999. When is a functional program not a functional program?. In *ICFP*. ACM, 1–7.

[29] John Longley. 2018. Bar recursion is not computable via iteration. To appear in Computability. Available at arxiv.org/abs/1804.07277. (2018).

[30] John Longley. 2018. The recursion hierarchy for PCF is strict. *Logical Methods in Comput. Sci.* 14, 3:8 (2018), 1–51.

[31] John Longley and Dag Normann. 2015. *Higher-Order Computability*. Springer.

[32] Chris Okasaki. 1999. *Purely functional data structures*. Cambridge University Press.

[33] Nicholas Pippenger. 1996. Pure versus impure Lisp. In *POPL*. ACM, 104–109.

[34] Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Typed Equivalence of Effect Handlers and Delimited Control. In *FSCD (LIPIcs)*, Vol. 131. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 30:1–30:16.

[35] Gordon Plotkin. 1997. LCF considered as a programming language. *Theor. Comput. Sci.* 5, 3 (1997), 223–255.

[36] Gordon D. Plotkin and John Power. 2001. Adequacy for Algebraic Effects. In *FoSSaCS (Lecture Notes in Computer Science)*, Vol. 2030. Springer, 1–24.

1431	[37] Gordon D. Plotkin and Matija Pretnar. 2013. Handling Algebraic Effects. <i>Logical Methods in Computer Science</i> 9, 4 (2013).	1486
1432		1487
1433	[38] Matija Pretnar. 2015. An Introduction to Algebraic Effects and Handlers. <i>Electr. Notes Theor. Comput. Sci.</i> 319 (2015), 19–35. Invited tutorial paper.	1488
1434		1489
1435	[39] Alex K. Simpson. 1998. Lazy Functional Algorithms for Exact Real Functionals. In <i>MFCS (Lecture Notes in Computer Science)</i> , Vol. 1450. Springer, 456–464.	1490
1436		1491
1437		1492
1438		1493
1439		1494
1440		1495
1441		1496
1442		1497
1443		1498
1444		1499
1445		1500
1446		1501
1447		1502
1448		1503
1449		1504
1450		1505
1451		1506
1452		1507
1453		1508
1454		1509
1455		1510
1456		1511
1457		1512
1458		1513
1459		1514
1460		1515
1461		1516
1462		1517
1463		1518
1464		1519
1465		1520
1466		1521
1467		1522
1468		1523
1469		1524
1470		1525
1471		1526
1472		1527
1473		1528
1474		1529
1475		1530
1476		1531
1477		1532
1478		1533
1479		1534
1480		1535
1481		1536
1482		1537
1483		1538
1484		1539
1485		1540

1541	Configurations	Continuations	1596
1542	$\langle\langle M \mid \gamma \mid \sigma \rangle\rangle = \langle\sigma\rangle(\langle M \rangle_\gamma)$	$\langle\langle \rangle\rangle M = M$	1597
1543		$\langle\langle \gamma, x, N \rangle\rangle :: \sigma \langle M \rangle = \langle\sigma\rangle(\mathbf{let} \ x \leftarrow M \ \mathbf{in} \ \langle N \rangle(\gamma \setminus \{x\}))$	1598
1544	Computation terms		1599
1545		$\langle\langle V \ W \rangle\rangle_\gamma = \langle V \rangle_\gamma \langle W \rangle_\gamma$	1600
1546		$\langle\langle \mathbf{let} \ \langle x; y \rangle = V \ \mathbf{in} \ N \rangle\rangle_\gamma = \mathbf{let} \ \langle x; y \rangle = \langle V \rangle_\gamma \ \mathbf{in} \ \langle N \rangle(\gamma \setminus \{x, y\})$	1601
1547		$\langle\langle \mathbf{case} \ V \ \{\mathbf{inl} \ x \mapsto M; \mathbf{inr} \ y \mapsto N\} \rangle\rangle_\gamma = \mathbf{case} \ \langle V \rangle_\gamma \ \{\mathbf{inl} \ x \mapsto \langle M \rangle(\gamma \setminus \{x\});$	1602
1548		$\mathbf{inr} \ y \mapsto \langle N \rangle(\gamma \setminus \{y\})\}$	1603
1549		$\langle\langle \mathbf{return} \ V \rangle\rangle_\gamma = \mathbf{return} \ \langle V \rangle_\gamma$	1604
1550		$\langle\langle \mathbf{let} \ x \leftarrow M \ \mathbf{in} \ N \rangle\rangle_\gamma = \mathbf{let} \ x \leftarrow \langle M \rangle_\gamma \ \mathbf{in} \ \langle N \rangle(\gamma \setminus \{x\})$	1605
1551	Value terms and values		1606
1552	$\langle\langle x \rangle\rangle_\gamma = \langle v \rangle$, if $\gamma(x) = v$	$\langle\langle n \rangle\rangle = n$	1607
1553	$\langle\langle x \rangle\rangle_\gamma = x$, if $x \notin \mathit{dom}(\gamma)$	$\langle\langle \gamma, \lambda x^A. M \rangle\rangle = \lambda x^A. \langle M \rangle(\gamma \setminus \{x\})$	1608
1554	$\langle\langle n \rangle\rangle_\gamma = n$	$\langle\langle \gamma, \mathbf{rec} \ f \ x^A. M \rangle\rangle = \mathbf{rec} \ f \ x^A. \langle M \rangle(\gamma \setminus \{f, x\})$	1609
1555	$\langle\langle \lambda x^A. M \rangle\rangle_\gamma = \lambda x^A. \langle M \rangle(\gamma \setminus \{x\})$	$\langle\langle \langle \rangle \rangle\rangle = \langle \rangle$	1610
1556	$\langle\langle \mathbf{rec} \ f \ x^A. M \rangle\rangle_\gamma = \mathbf{rec} \ f \ x^A. \langle M \rangle(\gamma \setminus \{f, x\})$	$\langle\langle \langle v; w \rangle \rangle\rangle = \langle \langle v \rangle \rangle; \langle \langle w \rangle \rangle$	1611
1557	$\langle\langle \langle \rangle \rangle\rangle_\gamma = \langle \rangle$	$\langle\langle \langle \mathbf{inl} \ v \rangle^B \rangle\rangle = \langle \mathbf{inl} \ \langle v \rangle \rangle^B$	1612
1558	$\langle\langle \langle V; W \rangle \rangle\rangle_\gamma = \langle \langle V \rangle \rangle_\gamma; \langle \langle W \rangle \rangle_\gamma$	$\langle\langle \langle \mathbf{inr} \ w \rangle^A \rangle\rangle = \langle \mathbf{inr} \ \langle w \rangle \rangle^A$	1613
1559	$\langle\langle \langle \mathbf{inl} \ V \rangle^B \rangle\rangle_\gamma = \langle \mathbf{inl} \ \langle V \rangle \rangle_\gamma^B$	$\langle\langle \langle \sigma^A \rangle \rangle\rangle = \lambda x^A. \langle\sigma\rangle(\mathbf{return} \ x)$	1614
1560	$\langle\langle \langle \mathbf{inr} \ W \rangle^A \rangle\rangle_\gamma = \langle \mathbf{inr} \ \langle W \rangle \rangle_\gamma^A$		1615

Figure 6. Mapping from Base Machine Configurations to Terms

A Proof Details for Correctness of the Base Abstract Machine

Correctness We now show that the abstract machine is correct with respect to the operational semantics, that is, the abstract machine faithfully simulates the operational semantics. Initial states provide a canonical way to map a computation term onto the abstract machine. A more interesting question is how to map an arbitrary configuration to a computation term. Figure 6 describes such a mapping $\langle\langle - \rangle\rangle$ from configurations to terms via a collection of mutually recursive functions defined on configurations, continuations, computation terms, value terms, and machine values. The mapping makes use of two operations on environments, γ , which we define now.

Definition A.1. We write $\mathit{dom}(\gamma)$ for the domain of γ , and $\gamma \setminus \{x_1, \dots, x_n\}$ for the restriction of environment γ to $\mathit{dom}(\gamma) \setminus \{x_1, \dots, x_n\}$.

The $\langle\langle - \rangle\rangle$ function enables us to classify the abstract machine reduction rules according to how they relate to the operational semantics. The rule (M-LET) is administrative in the sense that $\langle\langle - \rangle\rangle$ is invariant under this rule. This leaves the β -rules (M-APP), (M-SPLIT), (M-CASE), and (M-RETCONT). Each of these corresponds directly with performing a reduction in the operational semantics.

Definition A.2 (Auxiliary reduction relations). We write \longrightarrow_a for administrative steps, \longrightarrow_β for β -steps, and \Longrightarrow for a sequence of steps of the form $\longrightarrow_a^* \longrightarrow_\beta$.

The following lemma describes how we can simulate each reduction in the operational semantics by a sequence of administrative steps followed by one β -step in the abstract machine.

Lemma A.3. Suppose M is a computation and C is configuration such that $\langle\langle C \rangle\rangle = M$, then if $M \rightsquigarrow N$ there exists C' such that $C \Longrightarrow C'$ and $\langle\langle C' \rangle\rangle = N$, or if $M \not\rightsquigarrow$ then $C \not\Longrightarrow$.

Proof. By induction on the derivation of $M \rightsquigarrow N$. □

The correspondence here is rather strong: there is a one-to-one mapping between \rightsquigarrow and \Longrightarrow . The inverse of the lemma is straightforward as the semantics is deterministic. Notice that Lemma A.3 does not require that M be well-typed. We have chosen here not to perform type-erasure, but the results can be adapted to semantics in which all type annotations are erased.

Theorem A.4 (Simulation). If $\vdash M : A$ and $M \rightsquigarrow^+ N$ such that N is normal, then $\langle M \mid \emptyset \mid \langle \rangle \rangle \longrightarrow^+ C$ such that $\langle\langle C \rangle\rangle = N$, or $M \not\rightsquigarrow$ then $\langle M \mid \emptyset \mid \langle \rangle \rangle \not\longrightarrow$.

Proof. By repeated application of Lemma A.3. □

B Proof Details for the Complexity of Effectful Generic Counting

In this appendix we give proof details and artefacts for Theorem 5.12. Throughout this section we let H_{count} denote the handler definition of count, that is

$$H_{\text{count}} := \left\{ \begin{array}{l} \text{val } x \quad \mapsto \text{if } x \text{ then return 1 else return 0} \\ \text{Branch } \langle \rangle r \mapsto \text{let } x \leftarrow r \text{ true in} \\ \quad \text{let } y \leftarrow r \text{ false in} \\ \quad x + y \end{array} \right\}$$

The timed decision tree model embeds timing information. For the proof we must also know the abstract machine environment and the pure continuation. Thus we decorate timed decision trees with this information.

Definition B.1 (decorated timed decision trees). A decorated timed decision tree is a partial function $t : \mathbb{B}^* \rightarrow (\text{Lab} \times \text{Nat}) \times (\text{Env} \times \text{PureCont})$ such that its first projection $bs \mapsto t(bs).1$ is a timed decision tree. As an abbreviation, we define $\text{DT} := \mathbb{B}^* \rightarrow (\text{Lab} \times \text{Nat}) \times (\text{Env} \times \text{PureCont})$.

We extend the projections labs and steps in the obvious way to work over decorated timed decision trees. We define two further projections. The first $\text{env}(t) := bs \mapsto t(bs).2.1$ projects the environment, whilst the second $\text{pure}(t) := bs \mapsto t(bs).2.2$ projects the pure continuation.

The following definition gives a procedure for constructing a decorated timed decision tree. The construction is similar to that of Definition 5.4.

Definition B.2. The decorated timed decision tree of a configuration is defined by the following equations

$$\begin{aligned} \mathcal{D} &: \text{Conf} \rightarrow \text{DT} \\ \mathcal{D}(\langle \text{return true } | \gamma | [] \rangle) &= ((\text{!true}, 0), (\gamma, [])) \\ \mathcal{D}(\langle \text{return false } | \gamma | [] \rangle) &= ((\text{!false}, 0), (\gamma, [])) \\ \mathcal{D}(\langle p V | \gamma | \sigma \rangle) &= ((\text{?}[V]\gamma), 0), (\gamma, \sigma) \\ \mathcal{D}(\langle p V | \gamma | \sigma \rangle)(b :: bs) &\simeq \mathcal{D}(\langle \text{return } b | \gamma | \sigma \rangle) bs \\ \mathcal{D}(\langle M | \gamma | \sigma \rangle) bs &\simeq \mathcal{I}(\mathcal{D}(\langle M' | \gamma' | \sigma' \rangle) bs), \\ &\quad \text{if } \langle M | \gamma | \sigma \rangle \longrightarrow \langle M' | \gamma' | \sigma' \rangle \end{aligned}$$

where $\mathcal{I}((\ell, s), (\gamma, \sigma)) := ((\ell, s + 1), (\gamma, \sigma))$ and p is a distinguished free variable such that in all of the above equations $\gamma(p) = \gamma'(p) = p$.

We shall write $\mathcal{D}(P)$ to mean $\mathcal{D}(\langle P p | \emptyset[p \mapsto p] | [] \rangle)$.

We define some functions, that given a list of booleans and a n -standard predicate, compute configurations of the effectful abstract machine at particular points of interest during evaluation of the given predicate. Let $\chi_{\text{count}}(V) := (\emptyset[\text{pred} \mapsto \llbracket V \rrbracket \emptyset], H_{\text{count}})$ denote the handler closure of H_{count} .

Notation. For an n -standard predicate P we write $|P| = n$ for the size of the predicate. Furthermore, we define χ_{id} for the identity handler closure $(\emptyset, \{\text{val } x \mapsto x\})$.

Definition B.3 (computing machine configurations). For any given n -standard predicate P and a list of booleans bs , such that $|bs| \leq n$, we can compute machine configurations at points of interest during evaluation of count P .

To make the notation slightly simpler we use the following conventions whenever n , t , and c appear free: $n = |P|$, $t = \mathcal{D}(P)$, and $c = C(P)$.

- The function `arrive` either computes the configuration at a query node, if $|bs| < n$, or the configuration at an answer node.

$$\begin{aligned} \text{arrive} &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Conf} \\ \text{arrive}(bs, P) &:= \langle V j | \gamma | (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle, \quad \text{if } |bs| < n \\ \text{where } \gamma &= \text{env}(t)(bs), ?j = \text{labs}(t)(bs), \text{ and } \llbracket V \rrbracket \gamma = (\text{env}^\perp(P), \lambda_.\text{do Branch } \langle \rangle) \\ \text{arrive}(bs, P) &:= \langle \text{return } b | \gamma | ([], \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle, \quad \text{if } |bs| = n \\ \text{where } \gamma &= \text{env}(t)(bs) \text{ and } !b = \text{labs}(t)(bs) \end{aligned}$$

- Correspondingly, the depart function computes the configuration either after the completion of a query or handling of an answer.

$$\begin{aligned} \text{depart} &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Conf} \\ \text{depart}(bs, P) &:= \langle \text{return } m \mid \gamma \mid \text{residual}(bs, P) \rangle, \quad \text{if } |bs| < n \\ &\text{where } \gamma = \text{env}_{\text{false}}^{\uparrow}(bs, P) \text{ and } m = c(\text{true} :: bs) + c(\text{false} :: bs) \\ \text{depart}(bs, P) &:= \langle \text{return } m \mid \gamma \mid \text{residual}(bs, P) \rangle, \quad \text{if } |bs| = n \\ &\text{where } \gamma = \text{env}^{\downarrow}(P) \text{ and } m = c(bs) \end{aligned}$$

The two clauses of depart yield slightly different configurations. The first clause computes a configuration inside the operation clause of H_{count} . The configuration is exactly tail-configuration after summing up the two respective values returned by the two invocations of resumption. Whilst the second clause computes the tail-configuration inside of the success clause of H_{count} after handling a return value of the predicate.

- The residual function computes the residual continuation structure which contains the bits of computations to perform after handling a complete path in a decision tree.

$$\begin{aligned} \text{residual} &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Cont} \\ \text{residual}(bs, P) &:= [(\text{purecont}(bs, P), \chi_{id})] \end{aligned}$$

- The function purecont computes the pure continuation.

$$\begin{aligned} \text{purecont} &: \mathbb{B}^* \times \text{Val} \rightarrow \text{PureCont} \\ \text{purecont}([], P) &:= [] \\ \text{purecont}(\text{true} :: bs, P) &:= (\gamma, x_{\text{true}}, \text{let } x_{\text{false}} \leftarrow r \text{ false in } x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \\ &\text{where } \gamma = \text{env}_{\text{true}}^{\downarrow}(\text{true} :: bs, P) \\ \text{purecont}(\text{false} :: bs, P) &:= (\gamma, x_{\text{false}}, x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \\ &\text{where } \gamma = \text{env}_{\text{false}}^{\downarrow}(\text{false} :: bs, P) \end{aligned}$$

- The function env^{\downarrow} computes the initial environment of the handler. The family of functions $\text{env}_{b \in \mathbb{B}}^{\downarrow}$ contains two functions, one for each instantiation of b , which describe how to compute the environment prior *descending* down a branch as the result of invoking a resumption with b . Analogously, the functions in the family $\text{env}_{b \in \mathbb{B}}^{\uparrow}$ describe how to compute the environment after *ascending* from the resumptive exploration of a branch.

$$\begin{aligned} \text{env}^{\downarrow} &: \text{Val} \rightarrow \text{Env} \\ \text{env}^{\downarrow}(P) &:= \emptyset[\text{pred} \mapsto \llbracket P \rrbracket \emptyset] \end{aligned}$$

$$\begin{aligned} \text{env}_{\text{true}}^{\downarrow} &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Env} \\ \text{env}_{\text{true}}^{\downarrow}(bs, P) &:= \text{env}^{\downarrow}(V)[r \mapsto (\sigma, \chi_{\text{count}}(P))], \\ &\text{where } \sigma = \text{pure}(t)(bs) \end{aligned}$$

$$\begin{aligned} \text{env}_{\text{false}}^{\downarrow} &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Env} \\ \text{env}_{\text{false}}^{\downarrow}(bs, P) &:= \gamma[x \mapsto i], \\ &\text{where } \gamma = \text{env}_{\text{true}}^{\downarrow}(bs, P) \text{ and } i = c(\text{true} :: bs) \end{aligned}$$

$$\begin{aligned} \text{env}_{\text{false}}^{\uparrow} &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Env} \\ \text{env}_{\text{false}}^{\uparrow}(bs, P) &:= \gamma[y \mapsto j], \\ &\text{where } \gamma = \text{env}_{\text{false}}^{\downarrow}(bs, P) \text{ and } j = c(\text{false} :: bs) \end{aligned}$$

We require an auxiliary lemma, because we need to be able to reason about bits of predicate computation, specifically when the predicate is first applied, going from a departure configuration to an arrival configuration, and from a departure configuration to an answer configuration. The following lemma states that for an n -standard predicate, handler machine transitions in lock-step with the base machine.

For a given predicate P we write $\chi_{\text{count}}(P)^{\text{val}}$ to mean $\chi_{\text{count}}(P)^{\text{val}} = (\emptyset[\text{pred} \mapsto \llbracket P \rrbracket \emptyset], H_{\text{count}})^{\text{val}} = H_{\text{count}}^{\text{val}}$, that is the projection of the success clause of H_{count} .

Lemma B.4. *For any given n -standard predicate P and a list of booleans $bs \in \mathbb{B}^*$ such that $|bs| \leq n$ along with two value $V : \text{Bool}$ and $b \in \mathbb{B}$, then the base machine and handler machine transition in lock-step in either way*

1871 1. If $|bs| = []$, then

$$1872 \quad \begin{array}{l} \langle P \ p \mid \gamma \mid [] \rangle \\ \xrightarrow{\text{steps}(t)([])} \\ \langle p \ (i) \mid \gamma' \mid \sigma \rangle, \end{array}$$

1876 where $?i = \text{labs}(t)([])$, $\gamma = \emptyset[P \mapsto P]$, $\gamma' = \text{env}(t)([])$, and $\sigma = \text{pure}(t)([])$; implies the handler machine perform the same
1877 amount of transitions

$$1878 \quad \begin{array}{l} \langle P \ p \mid \gamma \mid ([], \chi_{\text{count}}(P)) :: \text{residual}(P, []) \rangle [(\lambda_ . \text{do Branch } \langle \rangle) / p] \\ \xrightarrow{\text{steps}(t)([])} \\ \langle p \ (i) \mid \gamma' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, []) \rangle [(\lambda_ . \text{do Branch } \langle \rangle) / p] \end{array}$$

1882 2. For $bs = b :: bs'$ we have the following two subcases

1883 • If $|bs| < n$, then

$$1884 \quad \begin{array}{l} \langle \text{return } b \mid \gamma \mid \sigma \rangle \\ \xrightarrow{\text{steps}(t)(b::bs)} \\ \langle p \ (i) \mid \gamma' \mid \sigma \rangle, \end{array}$$

1888 where $?i = \text{labs}(t)(b :: bs)$, $\gamma = \text{env}_b^\perp$, $\gamma' = \text{env}(t)(b :: bs)$, and $\sigma = \text{pure}(t)(bs)$; implies the handler machine perform the
1889 same amount of transitions

$$1890 \quad \begin{array}{l} \langle \text{return } (b) \mid \gamma \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c) \rangle [(\lambda_ . \text{do Branch } \langle \rangle) / p] \\ \xrightarrow{\text{steps}(t)(b::bs)} \\ \langle p \ (i) \mid \gamma' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c) \rangle [(\lambda_ . \text{do Branch } \langle \rangle) / p] \end{array}$$

1895 • If $|bs| = n$, then

$$1896 \quad \begin{array}{l} \langle \text{return } (b) \mid \gamma \mid \sigma \rangle \\ \xrightarrow{\text{steps}(t)(b::bs')} \\ \langle \text{return } (b') \mid \gamma' \mid [] \rangle, \end{array}$$

1899 where $!b' = \text{labs}(t)(b :: bs)$, $\gamma = \text{env}(t)(bs)$, $\gamma' = \text{env}(t)(b :: bs)$, and $\sigma = \text{pure}(t)(bs)$; implies the handler machine
1900 perform the same amount of transitions

$$1901 \quad \begin{array}{l} \langle \text{return } (b) \mid \gamma \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c) \rangle [(\lambda_ . \text{do Branch } \langle \rangle) / p] \\ \xrightarrow{\text{steps}(t)(b::bs')} \\ \langle \text{return } (b') \mid \gamma' \mid ([], \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c) \rangle [(\lambda_ . \text{do Branch } \langle \rangle) / p] \end{array}$$

1906 *Proof.* Proof by induction on the transition relation \longrightarrow . □

1907 Let $\text{control} : \text{Conf} \rightarrow \text{Val}$ denote a partial function that hoists a value out of a given machine configuration, that is

$$1908 \quad \text{control}(\langle M \mid \gamma \mid \kappa \rangle) := \begin{cases} \llbracket V \rrbracket \gamma & \text{if } M = \text{return } V \\ \perp & \text{otherwise} \end{cases}$$

1909 The following lemma performs most of the heavy lifting for the proof of Theorem 5.12.

1910 **Lemma B.5.** Suppose P is an n -standard predicate, then for any list of booleans $bs \in \mathbb{B}^*$ such that $|bs| \leq n$

$$1911 \quad \text{arrive}(bs, P) \rightsquigarrow^{T(bs, n)} \text{depart}(bs, P),$$

1912 and $\text{control}(\text{depart}(bs, P)) \leq 2^{n-|bs|}$ with the function T defined as

$$1913 \quad T(bs, n) = \begin{cases} 9 * (2^{n-|bs|} - 1) + 2^{n-|bs|+1} + \sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) & \text{if } |bs| < n \\ 2 & \text{if } |bs| = n \end{cases}$$

1914 *Proof.* By downward induction on bs .

Base step We have that $|bs| = n$. Since the predicate is n -standard we further have that $n \geq 1$. We proceed by direct calculation.

$$\begin{aligned}
& \text{arrive}(bs, P) \\
&= (\text{definition of arrive when } n = |bs|) \\
& \langle \mathbf{return} \ b \mid \gamma \mid (\llbracket _ \rrbracket, \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle \\
& \quad \text{where } \gamma = \text{env}(t)(bs) \text{ and } !b = \text{labs}(t)(bs) \\
&\longrightarrow (\text{M-RETHANDLER, } \chi_{\text{count}}(P)^{\text{val}} = \{\mathbf{val} \ x \mapsto \dots\}) \\
& \langle \mathbf{if} \ x \ \mathbf{then} \ \mathbf{return} \ 1 \ \mathbf{else} \ \mathbf{return} \ 0 \mid \gamma'[x \mapsto \llbracket b \rrbracket \gamma'] \mid \text{residual}(bs, P) \rangle \\
& \quad \text{where } \gamma' = \chi_{\text{count}}(P).1
\end{aligned}$$

The value b can assume either of two values. We consider first the case $b = \text{true}$.

$$\begin{aligned}
&= (\text{assumption } b = \text{true}, \text{ definition of } \llbracket _ \rrbracket \text{ (2 value steps)}) \\
& \langle \mathbf{if} \ x \ \mathbf{then} \ \mathbf{return} \ 1 \ \mathbf{else} \ \mathbf{return} \ 0 \mid \gamma'[x \mapsto \text{true}] \mid \text{residual}(bs, P) \rangle \\
&\longrightarrow (\text{M-IF-TT (and } \log |\gamma'[x \mapsto \text{true}]| = 1 \text{ environment operations)}) \\
& \langle \mathbf{return} \ 1 \mid \gamma'[x \mapsto \text{true}] \mid \text{residual}(bs, n, P, t, c) \rangle \\
&= (\text{definition of depart when } n = |bs|) \\
& \text{depart}(bs, P)
\end{aligned}$$

We have that $\text{control}(\text{depart}(bs, P)) = 1 \leq 2^0 = 2^{n-|bs|}$. Next, we consider the case when $b = \text{false}$.

$$\begin{aligned}
&= (\text{assumption } b = \text{false}, \text{ definition of } \llbracket _ \rrbracket \text{ (2 value steps)}) \\
& \langle \mathbf{if} \ x \ \mathbf{then} \ \mathbf{return} \ 1 \ \mathbf{else} \ \mathbf{return} \ 0 \mid \gamma'[x \mapsto \text{false}] \mid \text{residual}(bs, P) \rangle \\
&\longrightarrow (\text{M-IF-TT (and } \log |\gamma'[x \mapsto \text{false}]| = 1 \text{ environment operations)}) \\
& \langle \mathbf{return} \ 0 \mid \gamma'[x \mapsto \text{false}] \mid \text{residual}(bs, n, P, t, c) \rangle \\
&= (\text{definition of depart when } n = |bs|) \\
& \text{depart}(bs, P)
\end{aligned}$$

Again, we have that $\text{control}(\text{depart}(bs, P)) = 0 \leq 2^0 = 2^{n-|bs|}$.

Step analysis In either case, the machine uses exactly 2 transitions. Thus we get that

$$2 = T(bs, n), \quad \text{when } |bs| = n$$

Inductive step The induction hypothesis states that for all $b \in \mathbb{B}$ and $|bs| < n$

$$\text{arrive}(b :: bs, P) \rightsquigarrow^{T(b::bs, n)} \text{depart}(b :: bs, P),$$

such that $\text{control}(\text{depart}(b :: bs, P)) \leq 2^{n-|b::bs|}$. We proceed by direct calculation.

$$\begin{aligned}
& \text{arrive}(bs, P) \\
&= (\text{definition of arrive when } n < |bs|) \\
& \langle V \ j \mid \gamma \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle \\
& \quad \text{where } ?j = \text{labs}(t)(bs), \gamma = \text{env}(t)(bs), \sigma = \text{pure}(t)(bs), \text{ and } V = (\text{env}^\perp(P), \lambda_.\mathbf{do} \ \text{Branch} \ \langle \rangle) \\
&\longrightarrow (\text{M-APP}) \\
& \langle \mathbf{do} \ \text{Branch} \ \langle \rangle \mid \gamma'[_ \mapsto \llbracket j \rrbracket \gamma'] \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle \\
& \quad \text{where } \gamma' = \text{env}^\perp(P) \\
&\longrightarrow (\text{M-HANDLE-OP, } \chi_{\text{count}}(P)^{\text{Branch}} = \{\text{Branch} \ \langle \rangle \ r \mapsto \dots\}) \\
& \left\langle \begin{array}{l} \mathbf{let} \ x_{\text{true}} \leftarrow r \ \mathbf{true} \ \mathbf{in} \\ \mathbf{let} \ x_{\text{false}} \leftarrow r \ \mathbf{false} \ \mathbf{in} \ \mid \gamma[r \mapsto \llbracket (\sigma, \chi_{\text{count}}(P)) \rrbracket \gamma] \mid \text{residual}(bs, P) \end{array} \right\rangle \\
& \quad x_{\text{true}} + x_{\text{false}} \\
& \quad \text{where } \gamma = \text{env}^\perp(P) \\
&= (\text{definition of } \llbracket _ \rrbracket \text{ (1 value step)}) \\
& \left\langle \begin{array}{l} \mathbf{let} \ x_{\text{true}} \leftarrow r \ \mathbf{true} \ \mathbf{in} \\ \mathbf{let} \ x_{\text{false}} \leftarrow r \ \mathbf{false} \ \mathbf{in} \ \mid \gamma' \mid \text{residual}(bs, P) \end{array} \right\rangle \\
& \quad x_{\text{true}} + x_{\text{false}} \\
& \quad \text{where } \gamma' = \gamma[r \mapsto (\sigma, \chi_{\text{count}}(P))] \\
&\longrightarrow (\text{M-LET, definition of residual}) \\
& \langle r \ \text{true} \mid \gamma' \mid \text{residual}(\text{true} :: bs, P) \rangle \\
&\longrightarrow (\text{M-RESUME, } \llbracket r \rrbracket \gamma' = (\sigma, \chi_{\text{count}}(P)) \text{ (} \log |\gamma'| = 1 \text{ environment operations)}) \\
& \langle \mathbf{return} \ \text{true} \mid \gamma' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(\text{true} :: bs, P) \rangle
\end{aligned}$$

We now use Lemma B.4 to reason about the progress of the predicate computation σ . There are two cases consider, either $1 + |bs| < n$ or $1 + |bs| = n$.

Case 1 $1 + |bs| < n$. We obtain the following configuration.

$$\begin{aligned}
& \longrightarrow \text{steps}(t)(\text{true}::bs) \quad (\text{by Lemma B.4}) \\
& \quad \langle V j \mid \gamma'' \mid (\sigma', \chi_{\text{count}}(P)) :: \text{residual}(\text{true} :: bs, P) \rangle \\
& \quad \quad \text{where } ?j = \text{labs}(t)(\text{true} :: bs), \gamma'' = \text{env}(t)(\text{true} :: bs), \sigma' = \text{pure}(t)(\text{true} :: bs) \\
& \quad \quad \quad \text{and } \llbracket V \rrbracket \gamma'' = (\text{env}^\perp(P), \lambda_ \mathbf{do} \text{ Branch } \langle \rangle) \\
& = \quad (\text{definition of arrive when } 1 + |bs| < n) \\
& \quad \text{arrive}(\text{true} :: bs, P) \\
& \longrightarrow T(\text{true}::bs, n) \quad (\text{induction hypothesis}) \\
& \quad \text{depart}(\text{true} :: bs, P) \\
& = \quad (\text{definition of depart when } 1 + |bs| < n) \\
& \quad \langle \mathbf{return} \ i \mid \gamma \mid \text{residual}(\text{true} :: bs, P) \rangle \\
& \quad \quad \text{where } i = c(\text{true} :: \text{true} :: bs) + c(\text{false} :: \text{true} :: bs) \text{ and } \gamma = \text{env}_{\text{false}}^\uparrow(\text{true} :: bs, P) \\
& = \quad (\text{definition of residual and purecont}) \\
& \quad \langle \mathbf{return} \ i \mid \gamma \mid [(\gamma', x_{\text{true}}, \mathbf{let} \ x_{\text{false}} \leftarrow r \ \text{false} \ \mathbf{in} \ x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id}] \rangle \\
& \quad \quad \quad \text{where } \gamma' = \text{env}_{\text{true}}^\downarrow(bs, P) \\
& \longrightarrow \quad (\text{M-RETCONT}) \\
& \quad \langle \mathbf{let} \ x_{\text{false}} \leftarrow r \ \text{false} \ \mathbf{in} \ x_{\text{true}} + x_{\text{false}} \mid \gamma'' \mid [(\text{purecont}(bs, P), \chi_{id})] \rangle \\
& \quad \quad \quad \text{where } \gamma'' = \gamma'[x_{\text{true}} \mapsto \llbracket i \rrbracket \gamma'] \\
& \longrightarrow \quad (\text{M-LET}) \\
& \quad \langle r \ \text{false} \mid \gamma'' \mid [(\gamma'', x_{\text{false}}, x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id}] \rangle \\
& = \quad (\text{definition of purecont and residual}) \\
& \quad \langle r \ \text{false} \mid \gamma'' \mid \text{residual}(\text{false} :: bs, P) \rangle \\
& \longrightarrow \quad (\text{M-RESUME}) \\
& \quad \langle \mathbf{return} \ \text{false} \mid \gamma'' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(\text{false} :: bs, P) \rangle \\
& \quad \quad \quad \text{where } \sigma = \text{pure}(t)(bs) \\
& \longrightarrow \text{steps}(t)(\text{false}::bs) \quad (\text{by Lemma B.4 and assumption } |\text{false} :: bs| < n) \\
& \quad \langle V j \mid \gamma \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(\text{false} :: bs, P) \rangle \\
& \quad \text{where } ?j = \text{labs}(t)(\text{false} :: bs), \sigma = \text{pure}(t)(\text{false} :: bs), \gamma = \text{env}(t)(\text{false} :: bs) \\
& \quad \quad \text{and } \llbracket V \rrbracket \gamma = (\text{env}^\perp(P), \lambda_ \mathbf{do} \text{ Branch } \langle \rangle) \\
& = \quad (\text{definition of arrive when } 1 + |bs| < n) \\
& \quad \text{arrive}(\text{false} :: bs, P) \\
& \longrightarrow T(\text{false}::bs, n) \quad (\text{induction hypothesis}) \\
& \quad \text{depart}(\text{false} :: bs, P) \\
& = \quad (\text{definition of depart when } 1 + |bs| < n) \\
& \quad \langle \mathbf{return} \ j \mid \gamma \mid \text{residual}(\text{false} :: bs, P) \rangle \\
& \quad \quad \text{where } j = c(\text{true} :: \text{false} :: bs) + c(\text{false} :: \text{false} :: bs) \text{ and } \gamma = \text{env}_{\text{false}}^\uparrow(\text{false} :: bs, P) \\
& = \quad (\text{definition of residual and purecont}) \\
& \quad \langle \mathbf{return} \ j \mid \gamma \mid [(\gamma'', x_{\text{false}}, x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id}] \rangle \\
& \longrightarrow \quad (\text{M-RETCONT}) \\
& \quad \langle x_{\text{true}} + x_{\text{false}} \mid \gamma''[x_{\text{false}} \mapsto \llbracket j \rrbracket \gamma''] \mid \text{residual}(bs, P) \rangle \\
& \longrightarrow \quad (\text{M-PLUS}) \\
& \quad \langle \mathbf{return} \ m \mid \gamma''[x_{\text{false}} \mapsto \llbracket j \rrbracket \gamma''] \mid \text{residual}(bs, P) \rangle \\
& \quad \text{where } m = c(\text{true} :: \text{true} :: bs) + c(\text{false} :: \text{true} :: bs) + c(\text{true} :: \text{false} :: bs) + c(\text{false} :: \text{false} :: bs) \\
& \quad \quad = c(\text{true} :: bs) + c(\text{false} :: bs) = c(bs) \leq 2^{n-|bs|} \\
& = \quad (\text{definition of depart when } |bs| < n) \\
& \quad \text{depart}(bs, P)
\end{aligned}$$

Step analysis The total amount of machine transitions is given by

$$\begin{aligned}
& 9 + \text{steps}(t)(\text{true} :: bs) + T(\text{true} :: bs, n) + \text{steps}(t)(\text{false} :: bs) + T(\text{false} :: bs, n) \\
&= \text{(reorder)} \\
& 9 + T(\text{true} :: bs, n) + \text{steps}(t)(\text{false} :: bs) + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(definition of } T) \\
& 9 + 9 * (2^{n-|\text{true}::bs|} - 1) + 9 * (2^{n-|\text{false}::bs|} - 1) + 2^{n-|\text{true}::bs|+1} + 2^{n-|\text{false}::bs|+1} \\
& + \sum_{1 \leq |bs'| \leq n-|\text{true}::bs|} \text{steps}(t)(bs' ++ \text{true} :: bs) + \sum_{1 \leq |bs'| \leq n-|\text{false}::bs|} \text{steps}(t)(bs' ++ \text{false} :: bs) \\
& + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(simplify)} \\
& 9 + 9 * (2^{n-|\text{true}::bs|} - 1) + 9 * (2^{n-|\text{false}::bs|} - 1) + 2^{n-|bs|+1} \\
& + \sum_{1 \leq |bs'| \leq n-|\text{true}::bs|} \text{steps}(t)(bs' ++ \text{true} :: bs) + \sum_{1 \leq |bs'| \leq n-|\text{false}::bs|} \text{steps}(t)(bs' ++ \text{false} :: bs) \\
& + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(merge sums)} \\
& 9 + 9 * (2^{n-|\text{true}::bs|} - 1) + 9 * (2^{n-|\text{false}::bs|} - 1) + 2^{n-|bs|+1} \\
& + \left(\sum_{\substack{2 \leq |bs'| \leq n-|bs| \\ bs' \in \mathbb{B}^*}} \text{steps}(t)(bs' ++ bs) \right) + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(rewrite binary sum)} \\
& 9 + 9 * (2^{n-|\text{true}::bs|} - 1) + 9 * (2^{n-|\text{false}::bs|} - 1) + 2^{n-|bs|+1} \\
& + \sum_{2 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) + \sum_{1 \leq |bs'| \leq 1} \text{steps}(t)(bs' ++ bs) \\
&= \text{(merge sums)} \\
& 9 + 9 * (2^{n-|\text{true}::bs|} - 1) + 9 * (2^{n-|\text{false}::bs|} - 1) + 2^{n-|bs|+1} + \sum_{1 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) \\
&= \text{(factoring)} \\
& 9 + 2 * 9 * (2^{n-|bs|-1} - 1) + 2^{n-|bs|+1} + \sum_{1 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) \\
&= \text{(distribute)} \\
& 9 + 9 * (2^{n-|bs|} - 2) + 2^{n-|bs|+1} + \sum_{1 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) \\
&= \text{(distribute)} \\
& 9 + 9 * 2^{n-|bs|} - 18 + 2^{n-|bs|+1} + \sum_{1 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) \\
&= \text{(simplify)} \\
& 9 * 2^{n-|bs|} - 9 + 2^{n-|bs|+1} + \sum_{1 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) \\
&= \text{(factoring)} \\
& 9 * (2^{n-|bs|} - 1) + 2^{n-|bs|+1} + \sum_{1 \leq |bs'| \leq n-|bs|} \text{steps}(t)(bs' ++ bs) \\
&= \text{(definition of } T) \\
& T(bs, n)
\end{aligned}$$

Case $1 + |bs| = n$. We obtain the following configuration.

$$\begin{aligned}
&\longrightarrow \text{steps}(t)(\text{true}::bs) \quad (\text{by Lemma B.4}) \\
&\quad \langle \text{return } b \mid \gamma'' \mid ([], \chi_{\text{count}}(P)) :: \text{residual}(\text{true} :: bs, P) \rangle \\
&\quad \quad \text{where } !b = \text{labs}(t)(\text{true} :: bs), \gamma'' = \text{env}(t)(\text{true} :: bs) \\
&= (\text{definition of arrive when } 1 + |bs| = n) \\
&\quad \text{arrive}(\text{true} :: bs, P) \\
&\longrightarrow T(\text{true}::bs, n) \quad (\text{induction hypothesis}) \\
&\quad \text{depart}(\text{true} :: bs, P) \\
&= (\text{definition of depart when } 1 + |bs| = n) \\
&\quad \langle \text{return } i \mid \gamma \mid \text{residual}(\text{true} :: bs, P) \rangle \\
&\quad \quad \text{where } i = c(\text{true} :: bs) \leq 2^{n-|\text{true}::bs|} = 1 \text{ and } \gamma = \text{env}^\perp(P) \\
&= (\text{definition of residual and purecont}) \\
&\quad \langle \text{return } i \mid \gamma \mid [(\gamma', x_{\text{true}}, \text{let } x_{\text{false}} \leftarrow r \text{ false in } x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id}] \rangle \\
&\longrightarrow (\text{M-RETCONT}) \\
&\quad \langle \text{let } x_{\text{false}} \leftarrow r \text{ false in } x_{\text{true}} + x_{\text{false}} \mid \gamma'[x_{\text{true}} \mapsto [i]\gamma'] \mid [(\text{purecont}(bs, P), \chi_{id})] \rangle \\
&= (\text{definition of } \llbracket - \rrbracket \text{ (1 value step)}) \\
&\quad \langle \text{let } x_{\text{false}} \leftarrow r \text{ false in } x_{\text{true}} + x_{\text{false}} \mid \gamma'' \mid [(\text{purecont}(bs, P), \chi_{id})] \rangle \\
&\quad \quad \text{where } \gamma'' = \gamma'[x_{\text{true}} \mapsto i] \\
&\longrightarrow (\text{M-LET, definition of residual}) \\
&\quad \langle r \text{ false} \mid \gamma'' \mid \text{residual}(\text{false} :: bs, P) \rangle \\
&\longrightarrow (\text{M-RESUME}) \\
&\quad \langle \text{return } \text{false} \mid \gamma'' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(\text{false} :: bs, P) \rangle \\
&\quad \quad \text{where } \sigma = \text{pure}(t)(bs) \\
&\longrightarrow \text{steps}(t)(\text{false}::bs) \quad (\text{by Lemma B.4 and assumption } 1 + |bs| = n) \\
&\quad \langle \text{return } b \mid \gamma \mid ([], \chi_{\text{count}}(P)) :: \text{residual}(\text{false} :: bs, P) \rangle \\
&\quad \quad \text{where } !b = \text{labs}(t)(\text{false} :: bs), \gamma = \text{env}(t)(\text{false} :: bs) \\
&= (\text{definition of arrive when } 1 + |bs| = n) \\
&\quad \text{arrive}(\text{false} :: bs, P) \\
&\longrightarrow T(\text{false}::bs, n) \quad (\text{induction hypothesis}) \\
&\quad \text{depart}(\text{false} :: bs, P) \\
&= (\text{definition of depart when } 1 + |bs| = n) \\
&\quad \langle \text{return } j \mid \gamma \mid \text{residual}(\text{false} :: bs, P) \rangle \\
&\quad \quad \text{where } j = c(\text{false} :: bs) \leq 2^{n-|\text{false}::bs|} = 1 \text{ and } \gamma = \text{env}^\perp(P) \\
&= (\text{definition of residual and purecont}) \\
&\quad \langle \text{return } j \mid \gamma \mid [(\gamma', x_{\text{false}}, x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id}] \rangle \\
&\quad \quad \text{where } \gamma' = \text{env}_{\text{false}}^\downarrow(bs, P) \\
&\longrightarrow (\text{M-RETCONT}) \\
&\quad \langle x_{\text{true}} + x_{\text{false}} \mid \gamma'' \mid [(\text{purecont}(bs, P), \chi_{id})] \rangle \\
&\quad \quad \text{where } \gamma'' = \gamma'[x_{\text{false}} \mapsto [j]\gamma'] = \gamma'[x_{\text{false}} \mapsto j] \\
&\longrightarrow (\text{M-PLUS}) \\
&\quad \langle \text{return } m \mid \gamma'' \mid [(\text{purecont}(bs, P), \chi_{id})] \rangle \\
&\quad \quad \text{where } m = c(\text{true} :: bs) + c(\text{false} :: bs) \leq 2^{n-|bs|} \\
&= (\text{definition of residual and depart when } |bs| < n) \\
&\quad \text{depart}(bs, P)
\end{aligned}$$

Step analysis The total amount of machine transitions is given by

$$\begin{aligned}
& 9 + \text{steps}(t)(\text{true} :: bs) + T(\text{true} :: bs, n) + \text{steps}(t)(\text{false} :: bs) + T(\text{false} :: bs, n) \\
&= \text{(reorder)} \\
& 9 + T(\text{true} :: bs, n) + T(\text{false} :: bs, n) + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(definition of } T \text{ when } |bs| + 1 = n) \\
& 9 + 2 + 2 + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(simplify)} \\
& 9 + 2^2 + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(rewrite } 2 = n - |bs| + 1) \\
& 9 + 2^{n-|bs|+1} + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(multiply by 1)} \\
& 9 * (2^{n-|bs|} - 1) + 2^{n-|bs|+1} + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
&= \text{(rewrite binary sum)} \\
& 9 * (2^{n-|bs|} - 1) + 2^{n-|bs|} + \sum_{\substack{1 \leq |bs'| \leq n-|bs| \\ bs' \in \mathbb{B}^*}} \text{steps}(t)(bs' ++ bs) \\
&= \text{(definition of } T) \\
& T(bs, n)
\end{aligned}$$

□

The following theorem is a copy of Theorem 5.12.

Theorem B.6. For all $n > 0$ and any n -standard predicate P it holds that

1. The program `effcount` is a generic counting program, that is:

$$\text{effcount } P \rightsquigarrow^+ \text{ return } V, \text{ such that } \mathbb{N}[[V]] = C(P)([]) \leq 2^n$$

2. The runtime complexity of `effcount` P is given by the following formula:

$$\sum_{\substack{|bs| \leq n \\ bs \in \mathbb{B}^*}} \text{steps}(\mathcal{T}(P))(bs) + O(2^n)$$

Proof. The proof begins by direct calculation.

$$\begin{aligned}
& \langle \text{effcount } P \mid \emptyset \mid [([], \chi_{id})] \rangle && 2531 \\
& = \text{(definition of residual)} && 2532 \\
& \langle \text{effcount } P \mid \emptyset \mid \text{residual}(P, [], t, c) \rangle && 2533 \\
& \longrightarrow (\text{M-APP}, \llbracket \text{effcount} \rrbracket \emptyset = (\emptyset, \lambda \text{pred}. \dots)) && 2534 \\
& \langle \text{handle } \text{pred } (\lambda _ . \text{do Branch } \langle \rangle) \text{ with } H_{\text{count}} \mid \gamma \mid \text{residual}(P, []) \rangle && 2535 \\
& \text{where } \gamma = \text{env}^+(P) && 2536 \\
& \longrightarrow (\text{M-HANDLE}) && 2537 \\
& \langle \text{pred } (\lambda _ . \text{do Branch } \langle \rangle) \mid \gamma \mid ([], (\gamma, H_{\text{count}})) :: \text{residual}(P, []) \rangle && 2538 \\
& = \text{(definition of } \chi_{\text{count}}) && 2539 \\
& \langle \text{pred } (\lambda _ . \text{do Branch } \langle \rangle) \mid \gamma \mid ([], \chi_{\text{count}}(P)) :: \text{residual}(P, []) \rangle && 2540 \\
& \longrightarrow \text{steps}(t)([]) \text{ (by Lemma B.4)} && 2541 \\
& \langle (\lambda _ . \text{do Branch } \langle \rangle) j \mid \gamma' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, []) \rangle && 2542 \\
& \text{where } \gamma' = \text{env}(t)([]), \sigma = \text{pure}(t)(bs) \text{ and } ?j = \text{labs}(t)(bs) && 2543 \\
& = \text{(definition of arrive)} && 2544 \\
& \text{arrive}(P, []) && 2545 \\
& \longrightarrow T([], n) \text{ (by Lemma B.5)} && 2546 \\
& \text{depart}(P, []) && 2547 \\
& = \text{(definition of depart)} && 2548 \\
& \langle \text{return } m \mid \gamma \mid \text{residual}(P, []) \rangle && 2549 \\
& \text{where } \gamma = \text{env}^+(P) \text{ and } m = c([]) \leq 2^{n-|bs|} = 2^n && 2550 \\
& = \text{(definition of residual)} && 2551 \\
& \langle \text{return } m \mid \gamma \mid [([], \chi_{id})] \rangle && 2552 \\
& \longrightarrow (\text{M-HANDLE-RET}, H_{id}^{\text{val}} = \{\text{val } x \mapsto \text{return } x\}) && 2553 \\
& \langle \text{return } x \mid \emptyset[x \mapsto m] \mid [] \rangle && 2554
\end{aligned}$$

Analysis The machine yields the value m . By Lemma B.5 it follows that $m \leq 2^{n-|bs|} = 2^{n-|[]|} = 2^n$. Furthermore, the total amount of transitions used were

$$\begin{aligned}
& 5 + \text{steps}(t)([]) + T([], n) && 2555 \\
& = \text{(definition of } T) && 2556 \\
& 5 + \text{steps}(t)([]) + 9 * 2^n + 2^{n+1} + \sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') && 2557 \\
& = \text{(simplify)} && 2558 \\
& 5 + \text{steps}(t)([]) + 9 * 2^n + 2^{n+1} + \sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') && 2559 \\
& = \text{(reorder)} && 2560 \\
& 5 + \left(\sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') \right) + \text{steps}(t)([]) + 9 * 2^n + 2^{n+1} && 2561 \\
& = \text{(rewrite as unary sum)} && 2562 \\
& 5 + \left(\sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') + \sum_{bs' \in \mathbb{B}^*}^{0 \leq |bs'| \leq 0} \text{steps}(t)(bs') \right) + 9 * 2^n + 2^{n+1} && 2563 \\
& = \text{(merge sums)} && 2564 \\
& 5 + \left(\sum_{bs' \in \mathbb{B}^*}^{0 \leq |bs'| \leq n} \text{steps}(t)(bs') \right) + 9 * 2^n + 2^{n+1} && 2565 \\
& = \text{(definition of } \mathcal{O}) && 2566 \\
& \left(\sum_{bs' \in \mathbb{B}^*}^{0 \leq |bs'| \leq n} \text{steps}(t)(bs') \right) + \mathcal{O}(2^n) && 2567
\end{aligned}$$

□

C Proof Details for the No Shortcuts Lemma

The proof of Lemma 5.13 rely on the fact that any n -standard predicate has a canonical form. Section C.1 disseminate canonical predicates, whilst Section C.2 proves Lemma 5.13.

C.1 Canonical Predicates

The decision tree model (Definition 5.2) captures the interaction between a given predicate P and its point p . The interior nodes correspond to those places where P queries p , whilst the leaves represent answers ultimately conferred from the dialogue between the predicate and its point.

The abstract nature of the decision tree model means that concrete syntactic structure of the predicate is lost. Thus we cannot hope to reconstruct a particular predicate from its model. Indeed many syntactically distinct predicates may share the same model. However, we can construct *some* predicate from a given model, namely, the *canonical predicate*. Intuitively, the canonical predicate P' of P is a predicate which exhibits the same dialogue as P for every (valid) point.

Let $\mathcal{U}(P) := bs \mapsto \mathcal{T}(P)(bs)$.¹ denote the procedure for constructing an *untimed decision tree* of a given predicate P .

Definition C.1 (Canonical predicate). A canonical predicate P' of an n -standard predicate P is itself an n -standard predicate whose body (syntactically) consists entirely of **let**-bindings of point applications and whose continuation is either another **let**-expression of the same form or **return** b for some boolean b . Moreover, P' exhibits the same dialogue as P , that is for all $bs \in \mathbb{B}^*$ such that $|bs| \leq n$ that

$$\mathcal{U}(P)(bs) = \mathcal{U}(P')(bs)$$

Next we define a procedure for constructing canonical predicate of any given n -standard predicate.

Definition C.2 (Normalisation procedure for predicates). The meta-procedure `norm` takes as input an n -standard untimed decision tree, and outputs a program whose type is $\text{Point} \rightarrow \text{Bool}$, which is exactly the type of predicates. The procedure makes use of an auxiliary procedure `body` to generate the predicate body.

$$\begin{aligned} \text{norm} & : (\mathbb{B}^* \rightarrow \text{Lab}) \rightarrow \text{Val} \\ \text{norm}(t) & := \lambda p^{\text{Point}}. \text{body}(t, [], p) \\ \text{body} & : (\mathbb{B}^* \rightarrow \text{Lab}) \times \mathbb{B}^* \times \text{Val} \rightarrow \text{Comp} \\ \text{body}(t, bs, p) & := \begin{cases} \text{return } b & t(bs) = !b \\ \text{let } b \leftarrow p \text{ i in} \\ \text{if } b \text{ then } \text{body}(t, \text{true} :: bs, p) & \text{if } t(bs) = ?i \\ \text{else } \text{body}(t, \text{false} :: bs, p) \end{cases} \end{aligned}$$

As convenient notation we write $\text{norm}(P)$ to mean $\text{norm}(bs \mapsto \mathcal{U}(P)(bs))$. Next we show that the meta-procedure `norm` produces canonical predicates.

Lemma C.3. *Suppose P is an n -standard predicate then $P' := \text{norm}(P)$ is an n -standard predicate such that for all $bs \in \mathbb{B}^*$, $|bs| \leq n$*

$$\mathcal{U}(P)(bs) = \mathcal{U}(P')(bs')$$

Proof. By induction on n and `body`. □

Lemma C.4. *The procedure `norm` generates canonical predicates.*

Proof. First observe that the syntax produced by the `body` procedure of `norm` conforms with the syntactic restrictions of canonical predicates (Definition C.1). The rest follows as by Lemma C.3. □

C.2 No Shortcuts

We now have the necessary machinery to show that every n -count program in λ_b has at least exponential time complexity. The following lemma is a copy of Lemma 5.13.

Lemma C.5. *Let P be an n -standard predicate. Suppose C is an n -count program, then C must apply P to at least 2^n distinct n -points.*

Proof. Proof by contradiction. Pick a boolean sequence $bs \in \mathbb{B}^n$. Suppose there exists an n -count program C which does not construct the critical point p_c corresponding to bs . Let b be the answer yielded by P p . Now construct a predicate P' which yields the same answers as P except that at p_c it yields $\neg b$. Such a predicate can be constructed by negating b in the untimed decision tree model of P , i.e.

$$t' := bs' \mapsto \begin{cases} \neg b & \text{if } bs = bs' \\ \mathcal{U}(P)(bs') & \text{otherwise} \end{cases}$$

Then $P' = \text{norm}(t')$ constructs a canonical predicate, whose count is either one less or one more than that of P , that is at bs we have

$$|C(P')(bs) - C(P)(bs)| = 1$$

because $[] \sqsubset bs$ we further obtain that $|C(P')([]) - C(P)([])| = 1$. Now there are two cases to consider:

1. If $C P = C P'$ then C cannot be an n -count program, because $C(P)([]) \neq C(P')([])$, which contradicts the assumption.
2. If $C P \neq C P'$ then we continue to reason about the length of the reduction sequences arising from applications of P and P' .

Lemma C.6. *Let $\mathcal{F}[-]$ be any multi-hole context in C such that $\mathcal{F}[P] = C P$ and the type of $\mathcal{F}[P]$ is either Nat or Bool . If $\mathcal{F}[P] \rightsquigarrow^m \text{return } V$ then $\mathcal{F}[P'] \rightsquigarrow^* \text{return } V$ where the type of V is either Nat or Bool .*

Proof. Proof by induction on the length of the reduction sequence, m .

Base step We have that $m = 0$ which implies $\mathcal{F}[P] \rightsquigarrow^0 \text{return } V$ from which it follows that $\mathcal{F}[-]$ is simply $\text{return } V$, thus it follows immediately that $\mathcal{F}[P'] \rightsquigarrow^0 \text{return } V$.

Induction step We have that $m = 1 + m'$. The induction hypothesis is

$$\forall \mathcal{F}. \mathcal{F}[P] \rightsquigarrow^{m'} \text{return } V \quad \text{implies} \quad \mathcal{F}[P'] \rightsquigarrow^* \text{return } V.$$

There are two cases to consider depending on whether applications of P occur in \mathcal{F} .

Case $\mathcal{F}[P]$ is not an application of P . By assumption there is at least one reduction step, unroll this step to obtain

$$\mathcal{F}[-] \rightsquigarrow \mathcal{F}'[-] \rightsquigarrow^{m'} \text{return } V$$

Now plug in P' and then the result follows by a single application of the induction hypothesis.

Case $\mathcal{F}[P]$ is an application of P . It must be that P is applied to values of type Point . Moreover by assumption, we know that denotation of those values are distinct from the critical point p_c . Now write $\mathcal{F}[P] = \mathcal{G}[P, P p[P]]$ such that the first component of \mathcal{G} tracks residuals of P and the second component focuses on the expression in evaluation position, which in our particular case is an application of P to some point p in which P may occur again. We need to show that

$$\mathcal{G}[P, P p[P]] \rightsquigarrow \mathcal{G}[P, \text{return } W] \rightsquigarrow \text{return } V$$

for some $W : \text{Bool}$. Looking at the reduction sequence modulo $\mathcal{G}[P, -]$, we have that

$$P p[P] \rightsquigarrow^+ \mathcal{F}_0[p[P] i_0] \rightsquigarrow \mathcal{F}_0[\text{return } V_0] \rightsquigarrow^+ \mathcal{F}_1[p[P] i_1] \rightsquigarrow \dots \rightsquigarrow^+ \text{return } W,$$

where each reduction step is justified by the untimed decision tree model of P . From this we can deduce that

$$\mathcal{G}[P, P p[P]] \rightsquigarrow^+ \mathcal{G}[P, \text{return } W] \rightsquigarrow^* \text{return } V$$

where the last step follows by the induction hypothesis and $V : \text{Bool}$. Now, we argue that the above reduction sequence is tracked by $\mathcal{G}[P', -]$. The n -standardness of P' guarantees that it contains n queries, and moreover, since the decision tree model for P' is the same as P except for at one leaf, we know that the queries appear the in same order, so by appeal to the decision tree for P' we obtain that

$$P' p[P'] \rightsquigarrow^+ \mathcal{F}'_0[p[P'] i_0]$$

The term in evaluation position corresponds exactly to the first query node in the decision tree model. Now we can apply the induction hypothesis to obtain

$$\mathcal{F}'_0[p[P'] i_0] \rightsquigarrow^* \mathcal{F}'_0[\text{return } V_0]$$

The value V_0 is exactly the same answer to $p i_0$ as P obtained. Now there are two cases to consider depending on the value of n . If $n = 1$ then by the 1-standardness of P' we know that there will be no further queries, and it ultimately yields the same W as $P p$, because by assumption $p \neq p_c$. Otherwise if $n > 1$ then there must be further queries, and in particular, those queries must occur in the same order as those of P . Thus by the n -standardness of P' we get

$$\mathcal{F}'_0[\text{return } V_0] \rightsquigarrow^+ \mathcal{F}'_1[p[P'] i_1]$$

Yet again we find ourselves in a position where we can again apply the induction hypothesis to obtain an answer. By repeating this argument n times, we get that $P' p$ eventually yields W , we can lift this back into the outer context to obtain

$$\mathcal{G}[P', P' p[P']] \rightsquigarrow^+ \mathcal{G}[P', \mathbf{return} W]$$

and by the induction hypothesis, we get that

$$\mathcal{G}[P', \mathbf{return} W] \rightsquigarrow^* \mathbf{return} V.$$

□

Recall that $C P \neq C P'$, but by the Context Lemma C.6 both $C P$ and $C P'$ reduce to the same value which contradicts the initial assumption.

□

D Proof Details for the No Sharing Lemma

The following lemma is a copy of Lemma 5.18.

Lemma D.1. *Suppose P is an n -standard predicate and C is an n -count program, and let p_0 and p_1 be distinct n -points, then the predicate applications $P p_0$ and $P p_1$ within C have disjoint threads.*

Proof. Let $T_0 = \text{Th}(P p_0, \mathbb{P}[[p_0]])$ and $T_1 = \text{Th}(P p_1, \mathbb{P}[[p_1]])$ be the threads arising from the two distinct predicate applications. Suppose, without loss of generality, that P is applied to p_0 before p_1 , that is $C P \rightsquigarrow^+ \mathcal{E}_0[P p_0] \rightsquigarrow^+ \mathcal{E}_1[P p_1] \rightsquigarrow^+ \dots$ which by Lemma 5.16 implies that T_0 starts before T_1 . There are now two possible cases to consider.

1. T_0 finishes before T_1 starts. It follows immediately that T_0 and T_1 are disjoint.
2. T_1 starts in between the sections of T_0 . We now argue that T_1 must finish before evaluation of T_0 can continue. Suppose for any $i < n$ that the i -th query q_i starts T_1 , i.e

$$\mathcal{E}_i[p_0 q_i] \rightsquigarrow \mathcal{E}_i[\mathcal{E}'[P p_1]]$$

then by the ' n -ness' of C , P , and p_1 and since the reduction relation \rightsquigarrow is deterministic it follows that \mathcal{E}' reduces to a boolean value W which is plugged into the continuation of $\mathcal{E}_i[-]$

$$\mathcal{E}_i[p_0 q_i] \rightsquigarrow \mathcal{E}_i[\mathcal{E}'[P p_1]] \rightsquigarrow^+ \mathcal{E}_i[\mathbf{return} W]$$

Thus, T_1 must finish executing before evaluation of T_0 can resume.

□