

# Effects for Efficiency

Asymptotic Speedup with First-Class Control

DANIEL HILLERSTRÖM, The University of Edinburgh, UK

SAM LINDLEY, The University of Edinburgh and Imperial College London, UK

JOHN LONGLEY, The University of Edinburgh, UK

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 delimited control. 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  $\lambda_b$  and its extension with effect handlers  $\lambda_h$ . We show that  $\lambda_h$  admits an asymptotically more efficient implementation of generic search than any  $\lambda_b$  implementation of generic search. We also show that this efficiency gap remains when  $\lambda_b$  is extended with mutable state.

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, coroutines, 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 [1991] 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., [Knuth 1997, Section 2.6]), it is broadly taken as read that such models are equally

---

Authors’ addresses: 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.

50 applicable whatever programming language we are working with, and moreover that all respectable  
 51 languages can represent all algorithms of interest; thus, one does not expect the best achievable  
 52 asymptotic run-time for a typical algorithm (say in number theory or graph theory) to be sensitive  
 53 to the choice of programming language, except perhaps in marginal cases. (It should be admitted,  
 54 however, that proving general theorems to this effect may be harder than one might suppose: see  
 55 for example Section 1 of [Pippenger 1996].)

56 The situation changes radically, however, if we consider *higher-order* operations: programmable  
 57 operations whose inputs may themselves be programmable operations. (At this point, we suppose  
 58 that the languages we wish to compare all support higher-order data in some way: in particular,  
 59 that their type systems are rich enough to admit encodings of all simple types generated from the  
 60 familiar ground types via ‘ $\rightarrow$ ’.) Here it turns out that both what is computable and the efficiency  
 61 with which it can be computed can be highly sensitive to the selection of language features present.  
 62 This is in fact true more widely for *abstract data types*, of which higher-order types can be seen as  
 63 a special case: a higher-order value will of course be represented within the machine as ground  
 64 data, but a program within the language typically has no access to this internal representation, and  
 65 can interact with the value only by applying it to an argument.

66 Most work in this area to date has focused on computability differences. One of the best known  
 67 examples is the *parallel if* operation which is computable in a language with parallel evaluation  
 68 but not in a typical ‘sequential’ programming language [Plotkin 1977]. It is also well known that  
 69 the presence of control features or local state enables observational distinctions that cannot be  
 70 made in a purely functional setting: for instance, there are programs involving ‘call/cc’ that detect  
 71 the order in which a (call-by-name) ‘+’ operation evaluates its arguments [Cartwright and Felleisen  
 72 1992]. Such operations are ‘non-functional’ in the sense that their output is not determined solely  
 73 by the extension of their input (seen as a mathematical function  $\mathbb{N}_\perp \times \mathbb{N}_\perp \rightarrow \mathbb{N}_\perp$ ); however, there  
 74 are also programs with ‘functional’ behaviour that can be implemented with control or local state  
 75 but not without them [Longley 1999]. More recent results have exhibited differences lower down  
 76 in the language expressivity spectrum: for instance, in a purely functional setting *à la* Haskell, the  
 77 expressive power of *recursion* increases strictly with its type level [Longley 2018], and there are  
 78 natural operations computable by low-order recursion but not by high-order iteration [Longley  
 79 2019]. Much of this territory, including the mathematical theory of some of the natural notions of  
 80 higher-order computability that arise in this way, is mapped out by Longley and Normann [2015].

81 Relatively few results of this character have so far been established on the complexity side.  
 82 Pippenger [1996] gives an example of an ‘online’ operation on infinite sequences of atomic symbols  
 83 (essentially a function from streams to streams) such that the first  $n$  output symbols can be produced  
 84 within time  $O(n)$  if one is working in an ‘impure’ version of Lisp (in which mutation of ‘cons’ pairs  
 85 is admitted), but with a worst-case runtime no better than  $\Omega(n \log n)$  for any implementation  
 86 in pure Lisp (without such mutation). This example was reconsidered by Bird et al. [1997] who  
 87 showed that the same speedup can be achieved in a pure language by using lazy evaluation. Jones  
 88 [2001] explores the approach of manifesting expressivity and efficiency differences between certain  
 89 languages (which differ according to both the forms of iteration or recursion they admit and also  
 90 the use of higher types that they allow) by artificially restricting attention to ‘cons-free’ programs;  
 91 in this setting, the classes of representable first-order functions for the various languages are found  
 92 to coincide with some well-known complexity classes.

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

99 asking for an implementation of a certain third-order operation

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

101  
 102 A naive implementation strategy, supported by any reasonable language, is simply to apply  $P$  to  
 103 each of the  $2^n$  vectors in turn. A much less obvious, but still purely ‘functional’, approach due  
 104 to Berger [1990] achieves the effect of ‘pruned search’ where the predicate allows it (serving as  
 105 a warning that counter-intuitive phenomena can arise in this territory). Nonetheless, under a  
 106 mild condition on  $P$  (namely that it must inspect all  $n$  components of the given vector before  
 107 returning), both these approaches will have a  $\Omega(n2^n)$  runtime. Moreover, we shall show that in  
 108 a typical call-by-value language without advanced control features, one cannot improve on this:  
 109 *any* implementation of  $\text{count}_n$  must necessarily take time  $\Omega(n2^n)$ , even when the predicates  $P$  are  
 110 chosen to be ‘as simple as possible’. On the other hand, if we extend our language with a feature  
 111 such as *effect handlers* (see Section 2 below), it becomes possible to bring the runtime down to  
 112  $O(2^n)$ : an asymptotic gain of a factor of  $n$ .

113 The idea behind the speedup is easily explained and will already be familiar, at least informally,  
 114 to programmers who have worked with multi-shot continuations. Suppose for example  $n = 3$ , and  
 115 suppose that the predicate  $P$  always inspects the components of its argument in the order 0, 1, 2. A  
 116 naive implementation of  $\text{count}_3$  might start by applying the given  $P$  to  $p_0 = (\text{true}, \text{true}, \text{true})$ , and  
 117 then to  $p_1 = (\text{true}, \text{true}, \text{false})$ . Clearly there is some duplication here: the computations of  $P p_0$  and  
 118  $P p_1$  will proceed identically up to the point where the value of the final component is requested.  
 119 What we would like to do, then, is to record the state of the computation of  $P p_0$  at just this point,  
 120 so that we can later resume this computation with  $\text{false}$  supplied as the final component value in  
 121 order to obtain the value of  $P p_1$ . (Similarly for all other internal nodes in the evident binary tree  
 122 of boolean vectors.) Of course, this ‘backup’ approach would be standardly applied if one were  
 123 implementing a bespoke search operation for some *particular* choice of  $P$  (corresponding, say, to  
 124 the  $n$ -queens problem); but to apply this idea of resuming previous subcomputations in the generic  
 125 setting (that is, uniformly in  $P$ ) requires some special language feature such as effect handlers or  
 126 multi-shot continuations. One could also obviate the need for such a feature by choosing to present  
 127 the predicate  $P$  in some other way, but from our present perspective this would be to move the  
 128 goalposts: our intention is precisely to show that our languages differ in an essential way *as regards*  
 129 *their power to manipulate data of type*  $(\text{Nat} \rightarrow \text{Bool}) \rightarrow \text{Bool}$ .

130 This idea of using first-class control to achieve ‘backtracking’ has been exploited before and is  
 131 fairly widely known (see e.g. [Kiselyov et al. 2005]), and there is a clear programming intuition  
 132 that this yields a speedup unattainable in languages without such control features. Our main  
 133 contribution in this paper is to provide, for the first time, a precise mathematical theorem that pins  
 134 down this fundamental efficiency difference, thus giving formal substance to the above-mentioned  
 135 intuition. Since our goal is to give a realistic analysis of the efficiency achievable in various settings  
 136 without getting bogged down in inessential implementation details, we shall work concretely and  
 137 operationally with the languages in question, using a CEK-style abstract machine semantics as our  
 138 basic model of execution time, and with some specific programs in these languages. In the first  
 139 instance, we formulate our results as a comparison between a purely functional base language (a  
 140 version of call-by-value PCF) and an extension with first-class control; we then indicate how these  
 141 results can be extended to base languages with other features such as mutable state.

142 For their convenience as structured delimited control operators we adopt effect handlers as our  
 143 universal control abstraction of choice, but our results adapt mutatis mutandis to other first-class  
 144 control abstractions such as ‘call/cc’ [Sperber et al. 2009], ‘control’ ( $\mathcal{F}$ ) and ‘prompt’ ( $\#$ ) [Felleisen  
 145 1988], or ‘shift’ and ‘reset’ [Danvy and Filinski 1990].

146 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  $\lambda_b$  and its extension  $\lambda_h$  with effect handlers.
- Section 4 defines abstract machines for  $\lambda_b$  and  $\lambda_h$ , yielding a runtime cost model.
- Section 5 proves formal complexity results for generic search in  $\lambda_b$  ( $\Omega(n2^n)$ ) and  $\lambda_h$  ( $O(2^n)$ ).
- Section 6 shows that our results scale to richer settings including support for a wider class of predicates, an extension of the base language with state, and a non-trivial algorithm for generic search that exploits memoisation to perform pruned search.
- Section 7 evaluates implementations of generic search based on  $\lambda_b$  and  $\lambda_h$  in Standard ML.
- Section 8 concludes.

The languages  $\lambda_b$  and  $\lambda_h$  are rather minimal versions of previously studied systems – 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 [Plotkin and Power 2001; Plotkin and Pretnar 2013]. Subsequently they have emerged as a practical programming abstraction for modular effectful programming [Bauer and Pretnar 2015; Convent et al. 2020; Dolan et al. 2015; Hillerström et al. 2020; Kammar et al. 2013; Kiselyov et al. 2013; Leijen 2017]. 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 [2015], and as an introduction to the mathematical foundations of handlers, we refer the reader to the founding paper by Plotkin and Pretnar [2013] and the excellent tutorial paper by Bauer [2018].

Viewed through the lens of universal algebra, an algebraic effect is given by a signature  $\Sigma$  of finitary *operation symbols* defined over some nonempty carrier set  $A$ , along with an equational theory that describes the properties of the operations [Plotkin and Power 2001]. 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 a simple example, let us use the operation `Branch` to model a coin toss. Suppose we have a data type `Toss := Heads | Tails`, then we may implement a coin toss as follows.

```
toss : 1 → Toss
toss ⟨⟩ = if do Branch ⟨⟩ then Heads else 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 `⟨⟩` using the **do**-invocation form. The result of the operation determines whether the computation returns either `Heads` or `Tails`. Systems such as Frank [Convent et al. 2020; Lindley et al. 2017], Helium [Biernacki et al. 2019, 2020], Koka [Leijen 2017], and Links [Hillerström and Lindley 2016; Hillerström et al. 2020] include type-and-effect systems which track the use of effectful operations, whilst current iterations of systems such as Eff [Bauer and Pretnar 2015] and Multicore OCaml [Dolan et al. 2015] elect not to include an effect system. Our language is closer to the latter two.

We may, in the style of Lindley [2014], 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.

197	Types	$A, B, C, D ::= \text{Nat} \mid 1 \mid A \rightarrow B \mid A \times B \mid A + B$
198	Type Environments	$\Gamma ::= \cdot \mid \Gamma, x : A$
199	Values	$V, W \in \text{Val} ::= x \mid n \mid c \mid \lambda x^A. M \mid \mathbf{rec} f^A x. M$ $\mid \langle \rangle \mid \langle V, W \rangle \mid (\mathbf{inl} V)^B \mid (\mathbf{inr} W)^A$
200		
201	Computations	$M, N \in \text{Comp} ::= V W \mid \mathbf{let} \langle x, y \rangle = V \mathbf{in} N$ $\mid \mathbf{case} V \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \}$ $\mid \mathbf{return} V \mid \mathbf{let} x \leftarrow M \mathbf{in} N$
202		
203		
204		
205		
206		
207		
208		
209		
210		

Fig. 1. Syntax of  $\lambda_b$ 

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

```

214 handle toss  $\langle \rangle$  with  

215   val x  $\mapsto [x]$   

216   Branch  $\langle \rangle$   $r \mapsto r \text{ true} \# r \text{ false}$ 
  
```

217 The **handle**-construct generalises the exceptional syntax of Benton and Kennedy [2001]. A handler  
 218 has a *success* clause and an *operation* clause. The success clause determines how to interpret the  
 219 return value of toss, or equivalently how to interpret the leaves of its computation tree. It lifts the  
 220 return value into a singleton list. The operation clause determines how to interpret occurrences of  
 221 Branch in toss. It provides access to the argument of Branch (which is unit) and its resumption,  $r$ .  
 222 The resumption is a first-class delimited continuation which captures the remainder of the toss  
 223 computation from the invocation of Branch up to its nearest enclosing handler.

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

### 231 3 CALCULI

232 In this section, we present our base language  $\lambda_b$  and its extension with effect handlers  $\lambda_h$ .

#### 233 3.1 Base Calculus

234 The base calculus  $\lambda_b$  is a fine-grain call-by-value [Levy et al. 2003] variation of PCF [Plotkin 1977].  
 235 Fine-grain call-by-value is similar to A-normal form [Flanagan et al. 1993] in that every intermediate  
 236 computation is named, but unlike A-normal form is closed under reduction.

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

## Values

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

## Computations

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

Fig. 2. Typing Rules for  $\lambda_b$ 

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  $\text{Nat}$ , and  $r$  and  $k$  to denote resumptions and continuations, with the exception that we will use uppercase  $P$  to denote predicates. 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 will occasionally blur the distinction between object and meta language by writing  $A$  for the meta level type of closed value terms of type  $A$ . 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  $N$ .

The typing rules are given in Figure 2. 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  $\Gamma$ , where  $\square$  is either a value term ( $V$ ) or a computation term ( $M$ ). The constants have the following types.

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

We give a small-step operational semantics for  $\lambda_b$  with *evaluation contexts* in the style of Felleisen [1987]. The reduction rules are given in Figure 3. We write  $M[V/x]$  for  $M$  with  $V$  substituted for  $x$

295	S-APP	$(\lambda x^A. M)V \rightsquigarrow M[V/x]$	
296	S-APP-REC	$(\mathbf{rec} f^A x. M)V \rightsquigarrow M[(\mathbf{rec} f^A x. M)/f, V/x]$	
297	S-CONST	$c V \rightsquigarrow \mathbf{return} (\ulcorner c \urcorner (V))$	
298	S-SPLIT	$\mathbf{let} \langle x; y \rangle = \langle V; W \rangle \mathbf{in} N \rightsquigarrow N[V/x, W/y]$	
299	S-CASE-INL	$\mathbf{case} (\mathbf{inl} V)^B \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \} \rightsquigarrow M[V/x]$	
300	S-CASE-INR	$\mathbf{case} (\mathbf{inr} V)^A \{ \mathbf{inl} x \mapsto M; \mathbf{inr} y \mapsto N \} \rightsquigarrow N[V/y]$	
301	S-LET	$\mathbf{let} x \leftarrow \mathbf{return} V \mathbf{in} N \rightsquigarrow N[V/x]$	
302	S-LIFT	$\mathcal{E}[M] \rightsquigarrow \mathcal{E}[N],$	if $M \rightsquigarrow N$
303	Evaluation contexts $\mathcal{E} ::= [] \mid \mathbf{let} x \leftarrow \mathcal{E} \mathbf{in} N$		

Fig. 3. Contextual Small-Step Operational Semantics

and  $\ulcorner c \urcorner$  for the usual interpretation of constant  $c$  as a meta-level function on closed values. 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  $R$  and  $R^*$  for the reflexive, transitive closure of relation  $R$ . We write  $R/S$  for the quotient of relation  $R$  by relation  $S$ .

*Syntactic sugar.* 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 [Flanagan et al. 1993]. For example, the expression  $f(h w) + g \langle \rangle$  is syntactic sugar for:

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

We use the standard encoding of booleans as sums:

$$\begin{aligned} \text{Bool} &::= 1 + 1 & \text{true} &::= \mathbf{inl} \langle \rangle & \text{false} &::= \mathbf{inr} \langle \rangle \\ \mathbf{if} V \mathbf{then} M \mathbf{else} N &::= \mathbf{case} V \{ \mathbf{inl} \langle \rangle \mapsto M; \mathbf{inr} \langle \rangle \mapsto N \} \end{aligned}$$

We also define sequencing of computations in the standard way.

$$M; N ::= \mathbf{let} x \leftarrow M \mathbf{in} N, \quad \text{where } x \notin FV(N)$$

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.

### 3.2 Handler Calculus

We now define  $\lambda_h$  as an extension of  $\lambda_b$ . First we define notation for operation symbols, signatures, and handler types.

Operation symbols	$\ell \in \mathcal{L}$
Signatures	$\Sigma ::= \cdot \mid \{ \ell : A \rightarrow B \} \cup \Sigma$
Handler types	$F ::= C \Rightarrow D$

We assume a countably infinite set of operation symbols  $\mathcal{L}$ . An effect signature  $\Sigma$  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  $\Sigma$ . An effect handler type  $C \Rightarrow D$  classifies effect handlers that transform computations of type  $C$  into computations of type  $D$ . Following Pretnar [2015], we assume a global signature for every program.

### Computations

$$\begin{array}{c}
 \text{T-Do} \\
 (\ell : A \rightarrow B) \in \Sigma \quad \Gamma \vdash V : A \\
 \hline
 \Gamma \vdash \mathbf{do} \ell V : B \\
 \\
 \text{T-HANDLE} \\
 \Gamma \vdash M : C \quad \Gamma \vdash H : C \Rightarrow D \\
 \hline
 \Gamma \vdash \mathbf{handle} M \mathbf{with} H : D
 \end{array}$$

### Handlers

$$\begin{array}{c}
 \text{T-HANDLER} \\
 H^{\text{val}} = \{\mathbf{val} x \mapsto M\} \\
 [H^\ell = \{\ell p r \mapsto N_\ell\}]_{\ell \in \text{dom}(\Sigma)} \\
 [\Gamma, p : A_\ell, r : B_\ell \rightarrow D \vdash N_\ell : D]_{(\ell : A_\ell \rightarrow B_\ell) \in \Sigma} \\
 \Gamma, x : C \vdash M : D \\
 \hline
 \Gamma \vdash H : C \Rightarrow D
 \end{array}$$

Fig. 4. Additional Typing Rules for  $\lambda_h$

The typing rules for  $\lambda_h$  are those of  $\lambda_b$  (Figure 2) plus three additional rules for operations, handling, and handlers given in Figure 4. The T-Do rule ensures that an operation invocation is only well-typed if the operation  $\ell$  appears in the effect signature  $\Sigma$  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  $p$  ( $A_\ell$ ) and resumption  $r$  ( $B_\ell \rightarrow D$ ) in operation clause  $H^\ell$  is determined by the signature for  $\ell$ ; the return type of  $r$  is  $D$ , as the body of the resumption will itself be handled by  $H$ . We write  $H^\ell$  and  $H^{\text{val}}$  for projecting success and operation clauses.

$$\begin{array}{l}
 H^\ell := \{\ell p r \mapsto M\}, \quad \text{where } \{\ell p r \mapsto M\} \in H \\
 H^{\text{val}} := \{\mathbf{val} x \mapsto M\}, \quad \text{where } \{\mathbf{val} x \mapsto M\} \in H
 \end{array}$$

We extend the operational semantics to  $\lambda_h$ . Specifically, we add two new reduction rules: one for handling return values and another for handling operation invocations.

$$\begin{array}{l}
 \text{S-RET} \quad \mathbf{handle} (\mathbf{return} V) \mathbf{with} H \rightsquigarrow N[V/x], \quad \text{where } H^{\text{val}} = \{\mathbf{val} x \mapsto N\} \\
 \text{S-OP} \quad \mathbf{handle} \mathcal{E}[\mathbf{do} \ell V] \mathbf{with} H \rightsquigarrow N[V/p, \lambda y. \mathbf{handle} \mathcal{E}[\mathbf{return} y] \mathbf{with} H/r], \\
 \quad \text{where } H^\ell = \{\ell p r \mapsto N\}
 \end{array}$$

The first rule invokes the success clause. The second rule handles an operation via the corresponding operation clause. If we were to naively extend evaluation contexts with the handle construct then our semantics would become nondeterministic, as it may pick an arbitrary handlers in scope. In order to ensure that the semantics is deterministic, we instead add a distinct form of evaluation context for effectful computation, which we call handler contexts.

$$\text{Handler contexts } \mathcal{H} ::= [] \mid \mathbf{handle} \mathcal{H} \mathbf{with} H \mid \mathbf{let} x \leftarrow \mathcal{H} \mathbf{in} N$$

We replace the S-LIFT rule with a corresponding rule for handler contexts.

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

The separation between pure evaluation contexts  $\mathcal{E}$  and handler contexts  $\mathcal{H}$  ensures that the S-Op rule always selects the innermost handler.

We now characterise normal forms and state the standard type soundness property of  $\lambda_h$ .

*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  $\mathbf{return} V$ , or  $\mathcal{E}[\mathbf{do} \ell W]$ .

**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  $\lambda_b$  and its extension with effect handlers  $\lambda_h$ . 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 [Felleisen and Friedman 1987] for  $\lambda_b$  based on that of Hillerström et al. [2020]. 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 contains the environment  $\gamma$  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	$v, w \in \text{Mval} ::= x \mid n \mid c \mid \langle \rangle \mid \langle v, w \rangle$ $\mid (\gamma, \lambda x^A. M) \mid (\gamma, \mathbf{rec} f x^A. M) \mid (\mathbf{inl} v)^B \mid (\mathbf{inr} w)^A$
Pure continuations	$\sigma \in \text{PureCont} ::= [] \mid (\gamma, x, N) :: \sigma$

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

The abstract machine semantics is given in Figure 5. The transition relation ( $\longrightarrow$ ) makes use of the value interpretation ( $\llbracket - \rrbracket$ ) on value terms and machine values. The machine is initialised by placing a term in a configuration alongside the empty environment ( $\emptyset$ ) and identity pure 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 pure continuation with let bindings. The (M-RETCONT) rule extends the environment in the top frame of the pure continuation with a returned value. 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  $\gamma$ .

*Correctness.* The base machine faithfully simulates the operational semantics for  $\lambda_b$ ; most transitions correspond directly to  $\beta$ -reductions, but M-LET performs an administrative step to bring the computation  $M$  into evaluation position. We formally state and prove the correspondence in Appendix A, relying on an inverse map ( $\dashv$ ) from configurations to terms [Hillerström et al. 2020].

### 4.2 Handler Machine

We now enrich the  $\lambda_b$  machine to a  $\lambda_h$  machine. 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 in that the continuation component is replaced by a generalised continuation  $\kappa \in \text{Cont}$  [Hillerström et al. 2020]; a continuation is now a list of pairs containing a pure continuation (as in the base machine) and a handler closure ( $\chi$ ). A handler closure

**Transition relation**

442			
443	M-APP	$\langle V W \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma' [x \mapsto \llbracket W \rrbracket \gamma] \mid \sigma \rangle,$	
444			if $\llbracket V \rrbracket \gamma = (\gamma', \lambda x^A. M)$
445	M-REC	$\langle V W \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma' [f \mapsto (\gamma', \mathbf{rec} f x. M),$	
446		$x \mapsto \llbracket W \rrbracket \gamma] \mid \sigma \rangle,$	
447			if $\llbracket V \rrbracket \gamma = (\gamma', \mathbf{rec} f x^A. M)$
448	M-CONST	$\langle V W \mid \gamma \mid \sigma \rangle \longrightarrow \langle \mathbf{return} (\ulcorner c \urcorner (\llbracket V \rrbracket \gamma)) \mid \gamma \mid \sigma \rangle,$	
449			if $\llbracket V \rrbracket \gamma = c$
450	M-SPLIT	$\langle \mathbf{let} (x, y) = V \mathbf{in} N \mid \gamma \mid \sigma \rangle \longrightarrow \langle N \mid \gamma [x \mapsto v, y \mapsto w] \mid \sigma \rangle,$	
451			if $\llbracket V \rrbracket \gamma = \langle v; w \rangle$
452	M-CASEL	$\langle \mathbf{case} V \{ \mathbf{inl} x \mapsto M;$	
453		$\mathbf{inr} y \mapsto N \} \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma [x \mapsto v] \mid \sigma \rangle,$	
454			if $\llbracket V \rrbracket \gamma = \mathbf{inl} v$
455	M-CASER	$\langle \mathbf{case} V \{ \mathbf{inl} x \mapsto M;$	
456		$\mathbf{inr} y \mapsto N \} \mid \gamma \mid \sigma \rangle \longrightarrow \langle N \mid \gamma [y \mapsto v] \mid \sigma \rangle,$	
457			if $\llbracket V \rrbracket \gamma = \mathbf{inr} v$
458	M-LET	$\langle \mathbf{let} x \leftarrow M \mathbf{in} N \mid \gamma \mid \sigma \rangle \longrightarrow \langle M \mid \gamma \mid (\gamma, x, N) :: \sigma \rangle$	
459	M-RETCONT	$\langle \mathbf{return} V \mid \gamma \mid (\gamma', x, N) :: \sigma \rangle \longrightarrow \langle N \mid \gamma' [x \mapsto \llbracket V \rrbracket \gamma] \mid \sigma \rangle$	

**Value interpretation**

460	$\llbracket x \rrbracket \gamma = \gamma(x)$	$\llbracket n \rrbracket \gamma = n$	$\llbracket \lambda x^A. M \rrbracket \gamma = (\gamma, \lambda x^A. M)$
461	$\llbracket \langle \rangle \rrbracket \gamma = \langle \rangle$	$\llbracket c \rrbracket \gamma = c$	$\llbracket \mathbf{rec} f x^A. M \rrbracket \gamma = (\gamma, \mathbf{rec} f x^A. M)$
462			$\llbracket (\mathbf{inl} V)^B \rrbracket \gamma = (\mathbf{inl} \llbracket V \rrbracket \gamma)^B$
463	$\llbracket \langle V; W \rangle \rrbracket \gamma = \langle \llbracket V \rrbracket \gamma; \llbracket W \rrbracket \gamma \rangle$		$\llbracket (\mathbf{inr} V)^A \rrbracket \gamma = (\mathbf{inr} \llbracket V \rrbracket \gamma)^A$
464			

Fig. 5. Abstract Machine Semantics for  $\lambda_b$ **Transition relation**

467			
468	M-RESUME	$\langle V W \mid \gamma \mid \kappa \rangle \longrightarrow \langle \mathbf{return} W \mid \gamma \mid (\sigma, \chi) :: \kappa \rangle,$	
469			if $\llbracket V \rrbracket \gamma = (\sigma, \chi)^A$
470	M-LET	$\langle \mathbf{let} x \leftarrow M \mathbf{in} N \mid \gamma \mid (\sigma, \chi) :: \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ((\gamma, x, N) :: \sigma, \chi) :: \kappa \rangle$	
471	M-RETCONT	$\langle \mathbf{return} V \mid \gamma \mid ((\gamma', x, N) :: \sigma, \chi) :: \kappa \rangle \longrightarrow \langle N \mid \gamma' [x \mapsto \llbracket V \rrbracket \gamma] \mid (\sigma, \chi) :: \kappa \rangle$	
472	M-HANDLE	$\langle \mathbf{handle} M \mathbf{with} H \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ([], (\gamma, H)) :: \kappa \rangle$	
473	M-RETHANDLER	$\langle \mathbf{return} V \mid \gamma \mid ([], (\gamma', H)) :: \kappa \rangle \longrightarrow \langle M \mid \gamma' [x \mapsto \llbracket V \rrbracket \gamma] \mid \kappa \rangle,$	
474			if $H^{\mathbf{val}} = \{ \mathbf{val} x \mapsto M \}$
475	M-HANDLE-OP	$\langle \mathbf{do} \ell V \mid \gamma \mid (\sigma, (\gamma', H)) :: \kappa \rangle \longrightarrow \langle M \mid \gamma' [p \mapsto \llbracket V \rrbracket \gamma,$	
476		$r \mapsto (\sigma, (\gamma', H))] \mid \kappa \rangle,$	
477			if $\ell : A \rightarrow B \in \Sigma$
478			and $H^\ell = \{ \ell p r \mapsto M \}$
479			

Fig. 6. Abstract Machine Semantics for  $\lambda_h$ 

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:

$$\kappa_0 := ([], (\emptyset, \{ \mathbf{val} x \mapsto x \}))$$

Machine values are augmented to include 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).

491 The handler machine adds transition rules for handlers, and modifies (M-LET) and (M-RETCONT)  
 492 from the base machine to account for the richer continuation structure. Figure 6 depicts the new  
 493 and modified rules. The (M-HANDLE) rule pushes a handler closure along with an empty pure  
 494 continuation onto the continuation stack. The (M-RETHANDLER) rule transfers control to the  
 495 success clause of the current handler once the pure continuation is empty. The (M-HANDLE-OP)  
 496 rule transfers control to the matching operation clause on the topmost handler, and during the  
 497 process it reifies the handler closure. Finally, the (M-RESUME) rule applies a reified handler closure,  
 498 by pushing it onto the continuation stack. The handler machine has two possible final states: either  
 499 it yields a value or it gets stuck on an unhandled operation.

500  
 501 *Correctness.* The handler machine faithfully simulates the operational semantics of  $\lambda_h$ . Extending  
 502 the result for the base machine, we formally state and prove the correspondence in Appendix B.

### 503 4.3 Realisability and Asymptotic Complexity

504 As witnessed by the work of Hillerström and Lindley [2018] the machine structures are readily real-  
 505 isable using standard persistent functional data structures. Pure continuations on the base machine  
 506 and generalised continuations on the handler machine can be implemented using linked lists with a  
 507 time complexity of  $\mathcal{O}(1)$  for the extension operation ( $\_ :: \_$ ). The topmost pure continuation on the  
 508 handler machine may also be extended in time  $\mathcal{O}(1)$ , as extending it only requires reaching under  
 509 the topmost handler closure. Environments,  $\gamma$ , can be realised using a map, with a time complexity  
 510 of  $\mathcal{O}(\log |\gamma|)$  for extension and lookup [Okasaki 1999].

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

516  
 517 *Continuation copying.* On the handler machine the topmost continuation frame can be copied in  
 518 constant time due to the persistent runtime and the layout of machine continuations. An alternative  
 519 design would be to make the runtime non-persistent, as in MLton [2020], in which case copying a  
 520 continuation frame  $((\sigma, (\gamma, \_)) :: \_)$  would be a  $\mathcal{O}(|\sigma| + |\gamma|)$  time operation.

521  
 522 *Primitive operations on naturals.* Our model assumes that arithmetic operations on arbitrary  
 523 natural numbers take  $\mathcal{O}(1)$  time. This is common practice in the study of algorithms when the main  
 524 interest lies elsewhere (see [Cormen et al. 2009, Section 2.2]). If desired, one could adopt a more  
 525 refined cost model that accounted for the bit-level complexity of arithmetic operations; however,  
 526 doing so have essentially the same impact on both of the situations we are wishing to compare,  
 527 and thus would add nothing but noise to the overall analysis.

## 528 5 EFFICIENT GENERIC SEARCH

529  
 530 We now come to the crux of the paper. In this section we prove that  $\lambda_h$  accommodates some  
 531 programmable operations with an asymptotic runtime bound that cannot be achieved in  $\lambda_b$ . Whilst  
 532 the positive half of this claim essentially consolidates a known piece of folklore, the negative half  
 533 appears to be a genuinely new result. To obtain our results, it suffices to find just one efficient  
 534 program in  $\lambda_h$  and show that *no* equivalent program in  $\lambda_b$  can achieve the same asymptotic  
 535 complexity. We take *generic search* as our example.

536 Generic search is a modular search procedure that finds solutions to a given search problem  $P$ .  
 537 Generic search is agnostic to the specific instantiation of  $P$ , and as a result is applicable across a  
 538 wide spectrum of domains. Classic examples such as Sudoku solving [Bird 2006] and the  $n$ -queens

540 problem [Bell and Stevens 2009] can be cast as instances of generic search. Other instantiations  
 541 include problems from game theory such as computing Nash equilibria, problems from graph theory  
 542 such as graph colouring, and problems from real analysis such as real number integration [Daniels  
 543 2016; Simpson 1998].

544 To simplify the presentation, we compute the number of solutions (generic count), rather than  
 545 materialising all solutions (generic search). With little extra effort one can tweak the development to  
 546 compute exact solutions. Informally, a generic count program takes as input a predicate and returns  
 547 the number of times the predicate yields true. A predicate returns a boolean value which signifies  
 548 whether its input satisfies the predicate. As input a predicate takes a bit vector of length  $n > 0$ ,  
 549 which we represent as a first-order function  $\text{Nat} \rightarrow \text{Bool}$ . Ultimately we ask for implementations  
 550 of a program, count, whose type is

$$551 \quad \text{count}_n : ((\text{Nat}_n \rightarrow \text{Bool}) \rightarrow \text{Bool}) \rightarrow \text{Nat}$$

552 where  $\text{Nat}_n$  admits elements of the set  $\mathbb{N}_n := \{0, \dots, n - 1\}$ . We often omit the  $n$  index when  
 553 clear from context; in particular it does not appear explicitly in the types of our programs as our  
 554 formalism does not support dependent types.

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

## 558 5.1 Predicates and Points

559 Higher-order functions are the key to our modular formulation of generic search. We define a  
 560 predicate of size  $n$  as a closed value of the following type

$$561 \quad \text{Predicate}_n := \text{Point}_n \rightarrow \text{Bool}$$

562 where  $n$  is a natural number, and a point is also a closed value of the following type

$$563 \quad \text{Point}_n := \text{Nat}_n \rightarrow \text{Bool}$$

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

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

$$569 \quad p_2 := \lambda i. \text{if } i = 0 \text{ then true else if } i = 1 \text{ then false else } \perp i$$

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

571 Now let us move onto some example predicates. We can give a whole family of constant true  
 572 predicates. For example  $\text{tt}_0$  returns true irrespective of its point.

$$573 \quad \text{tt}_0 := \lambda p. \text{true}$$

574 We can define a variation,  $\text{tt}_2$ , which inspects two components of its point, but still returns true.

$$575 \quad \text{tt}_2 := \lambda p. p 1; p 0; \text{true}$$

576 This predicate is slightly more interesting than  $\text{tt}_0$  as it is defined only for points defined on  $\text{Nat}_n$   
 577 for  $n \geq 2$ . A predicate may inspect the same component of its point more than once

$$578 \quad \text{red}_1 := \lambda p. p 0; p 0$$

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

$$580 \quad \text{div}_1 := \text{rec } \text{div } p. \text{if } p 0 \text{ then } \text{div } p \text{ else false}$$

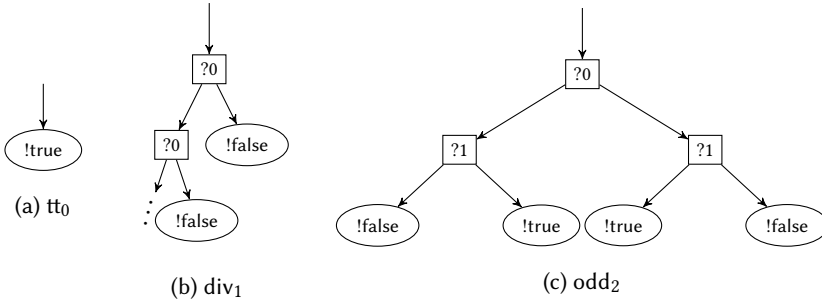


Fig. 7. Example Decision Tree Models

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 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  $\text{fold}$  and  $\text{map}$  are the standard combinators on lists and  $\otimes$  is exclusive-or. This predicate is only well-defined for  $n > 0$ . Applying  $\text{odd}_2$  to  $p_2$  yields true; applying it to  $p_{\text{true}}$  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 left subtree corresponds to  $p i$  being true and whose right subtree corresponds to  $p i$  being false, and each leaf is labelled with an answer !true or !false according to whether  $p$  satisfies the predicate. The trees are unrooted to account for the computation that occurs in between the application of a predicate to  $p$  and the first query or answer.

Figure 7 depicts models of some of the example predicates given above. The model of  $\text{tt}_0$  is simply an unrooted leaf (Figure 7a). The model of  $\text{div}_1$  is an infinite left-branching tree (Figure 7b). The model of  $\text{odd}_2$  is a complete binary tree (Figure 7c). A further example is the model of the unconditionally divergent predicate  $\text{div} := \text{rec } \text{div } p. \text{div } p$ , which is empty.

*Restrictions.* In order to obtain a meaningful complexity result we must constrain the predicates of interest. At one extreme, counting the size of a divergent predicate like  $\text{div}$  is meaningless. At the other extreme, a constant predicate like  $\text{tt}_0$  exhibits no interesting computational characteristics; other constant predicates like  $\text{tt}_2$  inspect their provided point. Predicates like  $\text{red}_1$  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 these conditions are  $\text{tt}_2$  and  $\text{odd}_n$ . Predicates satisfying 1 and 2 are exactly those whose models form complete binary trees (as in Figure 7c), 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  $\text{Nat}_n$ . 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 nondeterministic computation that we gave in Section 2. The main idea is to implement points as nondeterministic 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 as follows.

```

effcount : ((Nat → Bool) → Bool) → Nat
effcount P := handle P ( $\lambda\_.$ do Branch  $\langle \rangle$ ) with
  val b       $\mapsto$  if b then return 1 else return 0
  Branch  $\langle \rangle$  r  $\mapsto$  let xtrue  $\leftarrow$  r true in
    let xfalse  $\leftarrow$  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 ( $\lambda\_.$ **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  $\lambda_h$  than in  $\lambda_b$ .

## 5.3 Predicates, Points, and their Models, Formally

We now formalise the notions of  $n$ -standard predicates, points, and their models. For simplicity, we formalise these concepts using the operational semantics and abstract machine for the base language  $\lambda_b$ ; this means that the above concepts will be defined only for predicates expressible in  $\lambda_b$ . There is in principle no strong need for this restriction – with a little extra effort, corresponding concepts can be defined for  $\lambda_h$  predicates, and our efficiency result for effcount will be applicable to these too – but we choose to avoid this inessential complication.

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

$$\begin{aligned} \mathcal{T}(\langle \text{return true} \mid \gamma \mid [] \rangle) &= (\text{!true}, 0) & \mathcal{T}(\langle p V \mid \gamma \mid \sigma \rangle)(b :: bs) &\simeq \mathcal{T}(\langle \text{return } b \mid \gamma \mid \sigma \rangle) bs \\ \mathcal{T}(\langle \text{return false} \mid \gamma \mid [] \rangle) &= (\text{!false}, 0) & \mathcal{T}(\langle M \mid \gamma \mid \sigma \rangle) bs &\simeq \mathcal{I}(\mathcal{T}(\langle M' \mid \gamma' \mid \sigma' \rangle) bs), \\ \mathcal{T}(\langle p V \mid \gamma \mid \sigma \rangle) &= (? \llbracket V \rrbracket \gamma, 0) & \text{if } \langle M \mid \gamma \mid \sigma \rangle &\longrightarrow \langle M' \mid \gamma' \mid \sigma' \rangle \end{aligned}$$

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:  $\mathcal{T}(M) := \mathcal{T}(\langle M, \emptyset[p \mapsto p], \kappa_0 \rangle)$ . The decision tree of a predicate  $P$  is  $\mathcal{T}(P p)$ . Since  $p$  is a distinguished variable, we often omit it and write  $\mathcal{T}(P)$  for  $\mathcal{T}(P p)$ .

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  $b \in \mathbb{B}$ ; and
- 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  $t(bs) = t(bs') = ?j$  then  $bs = bs'$ .

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  $\lambda_b$ -language as we can reason intensionally about the single point used by  $\text{effcount}$  in the  $\lambda_h$ -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 ( $\text{Nat}_n$  for some  $n > 0$ ).

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

$$\forall i \in \mathbb{N}_n. p i \rightsquigarrow^* \text{return true} \vee p i \rightsquigarrow^* \text{return false}.$$

A semantic  $n$ -point  $\pi$  is the denotation of an  $n$ -point  $p$ , i.e. a mathematical function  $\mathbb{N}_n \rightarrow \text{Bool}$ . For any  $n$ -point  $p$  its corresponding semantic  $n$ -point is given by  $\pi = \mathbb{P}[\![p]\!]$ , where  $\mathbb{P}[\![-]\!]$  is the realisation of the operational behaviour of  $p$

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

Moreover, any two  $n$ -points  $p_0$  and  $p_1$  are said to be *distinct* if their corresponding semantic  $n$ -points differ, i.e.:

$$\exists i \in \mathbb{N}_n. \mathbb{P}[[P_0]] i \neq \mathbb{P}[[P_1]] i$$

## 5.4 Specification of Generic Counting

We now formally define generic counting.

*Definition 5.7.* 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 5.8.* The counting function for a configuration is defined by the following equations.

$$C(\langle \text{return true} \mid \gamma \mid [] \rangle) [] = 1$$

$$C(\langle \text{return false} \mid \gamma \mid [] \rangle) [] = 0$$

$$C(\langle p V \mid \gamma \mid \sigma \rangle) [] = C(\langle \text{return true} \mid \gamma \mid \sigma \rangle) [] + C(\langle \text{return false} \mid \gamma \mid \sigma \rangle) []$$

$$C(\langle p V \mid \gamma \mid \sigma \rangle) (b :: bs) \simeq C(\langle \text{return } b \mid \gamma \mid \sigma \rangle) bs$$

$$C(\langle M \mid \gamma \mid \sigma \rangle) bs \simeq C(\langle M' \mid \gamma' \mid \sigma' \rangle) bs, \quad \text{if } \langle M \mid \gamma \mid \sigma \rangle \longrightarrow \langle M' \mid \gamma' \mid \sigma' \rangle$$

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 5.9 (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  $P$

$$C P \rightsquigarrow^+ \text{return } C(P)([])$$

The restriction to  $n$ -standard predicates might at first seem rather tiresome and unnatural, but in the context of our work it has two motivations. First, it allows us to present the essence of our effectful generic counting algorithm in its simplest, cleanest form (compare the `effcount` program given above with the more widely applicable versions in Section 6.1 below). Second, it will enable us in Section 5.6 to present our main negative result in a particularly sharp form: in the base language  $\lambda_b$ , no  $n$ -count program can compete with `effcount` *even on  $n$ -standard predicates*.

## 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 C.

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

- (1) *The program `effcount` is a generic counting program.*
- (2) *The runtime complexity of `effcount`  $P$  is given by:*

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



PROOF. Both items can be proved by downwards induction on the length of  $bs$  and alternating, as needed, between intensional reasoning about reduction steps within `effcount` and extensional reasoning about reduction steps for  $P$ . We give the full details in Appendix C.

□

The above formula can clearly be simplified for certain reasonable classes of predicates. For instance, suppose we fix some constant  $c \in \mathbb{N}$ , and let  $\mathcal{P}_{n,c}$  be the class of all  $n$ -standard predicates  $P$  for which all the edge times  $\text{steps}(\mathcal{T}(P))(bs)$  are bounded by  $c$ . (Clearly, many reasonable predicates will belong to  $\mathcal{P}_{n,c}$  for some modest value of  $c$ .) Since the number of sequences  $bs$  in question is less than  $2^{n+1}$ , we may read off from the above formula that for predicates in  $\mathcal{P}_{n,c}$ , the runtime complexity of `effcount` is  $O(2^n)$ .

As a related aside, one might also ask about the execution time for an implementation of  $\lambda_h$  that performs genuine copying of continuations, as in systems such as MLton [2020]. We will not present the details of such an implementation, but it is informally clear that our  $O(2^n)$  bound would still apply as long as the continuations associated with internal nodes of  $\mathcal{T}(P)$  never becomes too large. Specifically, we might consider a class  $\mathcal{Q}_{n,c,k}$  of  $n$ -standard predicates  $P$  for which the edge times in  $\mathcal{T}(P)$  never exceed  $c$  and the sizes of the continuations never exceed  $k$ . (Once again, for reasonable  $c$  and  $k$  this gives us a respectable class of predicates.) Then it is intuitively clear that for such predicates, the total continuation-copying time will be  $O(2^n)$ , so that the overall runtime will still be  $O(2^n)$ .

## 5.6 Pure Generic Counting

We have shown that there is an implementation of `count` in  $\lambda_h$  with a runtime bound of  $O(2^n)$  for certain well-behaved predicates. We now prove that no implementation of `count` in  $\lambda_b$  can match this: in fact, we establish a *lower* bound of  $\Omega(n2^n)$  for the runtime of `count` on *any*  $n$ -standard predicate. Later, we shall extend our result to richer languages incorporating state or exceptions. This mathematically rigorous characterisation of the efficiency gap between languages with and without first-class control constructs is the central contribution of the paper.

One might ask at this point whether the claimed lower bound could not be obviated by means of some known continuation passing style (CPS) or monadic transform of effect handlers [Hillerström et al. 2017; Leijen 2017]. This can indeed be done, but only by dint of changing the type of our predicates  $P$  – which, as noted in the introduction, would defeat the purpose of our present enquiry. Our intention is precisely to investigate the relative power of various languages for manipulating predicates that are presented to us in a certain way which we do not have the luxury of choosing.

To get a feel for the issues that our proof must address, let us consider how one might go about constructing a `count` program in  $\lambda_b$ . The naive approach, of course, would be simply to apply the given predicate  $P$  to all  $2^n$  possible  $n$ -points in turn, keeping a count of those on which  $P$  yields true. It is a routine exercise to implement this approach in  $\lambda_b$ , yielding (parametrically in  $n$ ) a program

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

Since the evaluation of an  $n$ -standard predicate on an individual  $n$ -point  $p$  must clearly take time  $\Omega(n)$ , we have that the evaluation of `naivecountn` on any  $n$ -standard predicate  $P$  must take time  $\Omega(n2^n)$ . If  $P$  is not  $n$ -standard, the  $\Omega(n)$  lower bound need not apply, but we may still say that the evaluation of `naivecountn` on *any* predicate  $P$  (at level  $n$ ) must take time  $\Omega(2^n)$ .

One might at first suppose that both these properties are inevitable for any implementation of `count` within  $\lambda_b$ , or indeed any purely functional language: surely, the only way to learn something about the behaviour of  $P$  on every possible  $n$ -point is to apply  $P$  to each of these points in turn? It turns out, however, that the  $\Omega(2^n)$  lower bound can sometimes be circumvented by implementations that cleverly exploit *nesting* of calls to  $P$ . The germ of the idea may be illustrated within  $\lambda_b$  itself.

834 Suppose that we first construct some program

835  $\text{bestshot}_n : ((\text{Nat}_n \rightarrow \text{Bool}) \rightarrow \text{Bool}) \rightarrow (\text{Nat}_n \rightarrow \text{Bool})$   
 836

837 which, given a predicate  $P$ , returns some  $n$ -point  $p$  such that  $P p$  evaluates to true whenever this is  
 838 possible (i.e. whenever some such point exists). If  $P$  returns false on every  $n$ -point, we require simply  
 839 that  $\text{bestshot}_n P$  returns some arbitrary  $n$ -point. (In other words,  $\text{bestshot}_n$  embodies Hilbert's  
 840 choice operator  $\varepsilon$  on predicates.) It is once again routine to construct such a program by naive  
 841 means; and we may moreover assume that for any  $P$ , the evaluation of  $\text{bestshot}_n P$  takes only  
 842 constant time, all the real work being deferred until the argument of type  $\text{Nat}_n$  is supplied.

843 Now consider the following program:

844  $\text{lazycount}_n := \lambda P. \text{if } P (\text{bestshot}_n P) \text{ then naivecount}_n P \text{ else return } 0$   
 845

846 Here the term  $P (\text{bestshot}_n P)$  serves to test whether there exists an  $n$ -point satisfying  $P$ : if there  
 847 is not, our count program may return 0 straightaway. It is thus clear that  $\text{lazycount}_n$  is a correct  
 848 implementation of generic counting, and also that if  $P$  is the predicate  $\lambda p. \text{false}$  then  $\text{lazycount}_n P$   
 849 will return 0 within  $O(1)$  time, thus violating the  $\Omega(2^n)$  lower bound suggested above.

850 This might seem a rather footling point, as  $\text{lazycount}_n$  offers this efficiency gain *only* on (some  
 851 implementations of) the everywhere false predicate. However, by means of a recursive application  
 852 of such a nesting trick, we may arrive at a generic count program that spectacularly defies the  
 853  $\Omega(2^n)$  lower bound for an interesting class of (non- $n$ -standard) predicates, and indeed proves quite  
 854 viable for counting solutions to ' $n$ -queens' and similar problems. We shall refer to this program  
 855 *BergerCount*, since it is modelled largely on Berger's PCF implementation of the so-called *fan*  
 856 *functional* ([Berger 1990]; see also [Longley and Normann 2015]). This program is of some interest  
 857 in its own right, and will be briefly presented in Section 6.3. As we shall see, *BergerCount* actually  
 858 requires a mild extension of  $\lambda_b$  with a 'memoisation' primitive to achieve the effect of call-by-need  
 859 evaluation; but such a language can still be seen as purely 'functional' in the same sense as Haskell.

860 In the meantime, however, the moral is that the use of *nesting* can lead to surprising phenomena  
 861 which sometimes defy intuition (Escardó [2007] gives some striking further examples of this). What  
 862 we now wish to show is that for *n-standard* predicates, the naive lower bound of  $\Omega(n2^n)$  cannot in  
 863 fact be circumvented; the example of *BergerCount* both highlights the need for a rigorous proof of  
 864 this and tells us that our argument will need to pay particular attention to the possibility of nesting.

865 We now proceed to the proof itself. In the interests of clarity, we first present a proof in the basic  
 866 setting of  $\lambda_b$ ; later we will see how the approach scales to languages with state (Section 6.2).

867 As a modest first step, we note that where lower bounds are concerned, it will suffice to work with  
 868 the small-step operational semantics of  $\lambda_b$  rather than the more elaborate abstract machine model  
 869 employed in Section 4.1. This is because, as observed in Section 4.1, there is a tight correspondence  
 870 between these two execution models such that for the evaluation of any closed term, the number of  
 871 abstract machine steps is always at least the number of small-step reductions. Thus, if we are able  
 872 to show that the number of small-step reductions for any count program in  $\lambda_b$  on any  $n$ -standard  
 873 predicate is  $\Omega(n2^n)$ , this will establish the desired lower bound on the runtime.

874 We now establish a key lemma, which vindicates the naive intuition that in the  $n$ -standard case,  
 875 the only way to discover the correct value for  $\text{count } P$  is to perform  $2^n$  separate applications  $P p$   
 876 (albeit allowing for the possibility that these applications need not be performed 'in turn' but might  
 877 be nested in some complex way). We outline the proof here; full details are in Appendix D.

878 **LEMMA 5.11 (NO SHORTCUTS).** *If  $C$  is an  $n$ -count program and  $P$  is an  $n$ -standard predicate, then  $C$*   
 879 *applies  $P$  to at least  $2^n$  distinct  $n$ -points. More formally, for any of the  $2^n$  possible semantic  $n$ -points*  
 880  *$\pi : \mathbb{N}_n \rightarrow \mathbb{B}$ , there is a term  $\mathcal{E}[P p]$  appearing in the small-step reduction of  $C P$  such that  $p$  is a closed*  
 881 *value (hence an  $n$ -point) and  $\mathbb{P}[\![p]\!] = \pi$ .*  
 882

PROOF. Suppose  $C$  and  $P$  are as above, and suppose for contradiction that  $\pi$  is some semantic  $n$ -point such that no corresponding application  $P p$  ever arises in the course of computing  $C P$ . Let  $t$  be the untimed decision tree for  $P$ . Now consider the leaf node in  $t$  corresponding to the point  $\pi$ , and let  $t'$  be the tree obtained from  $t$  by simply negating the boolean value at this leaf node. It is then a fairly simple matter to construct a predicate  $P'$  whose decision tree is  $t'$ .

Since the numbers of true-leaves in  $t$  and  $t'$  differ by 1, it is clear that if  $C$  is indeed a correct  $n$ -count program, then the values returned by  $C P$  and  $C P'$  will have an absolute difference of 1. On the other hand, we will argue that if the computation of  $C P$  never actually ‘visits’ the leaf node in question, then  $C$  is unable to detect any difference between  $P$  and  $P'$ .

The situation here is reminiscent of Milner’s *context lemma* for PCF [Milner 1977], which (loosely) says that essentially the only way to observe a difference between two programs is to apply them to some argument on which they differ. Traditional proofs of the context lemma reason by induction on length of reduction sequences, and our present proof is modelled on these. Specifically, one proves the following by induction on  $m$ :

Suppose  $C P \rightsquigarrow^* \mathcal{E}[P p[P]]$  where  $\mathcal{E}$  is an evaluation context, and the context  $p[-]$  abstracts all occurrences of  $P$  that are residuals of the key occurrence in  $C P$ . If  $P p[P] \rightsquigarrow^m \mathbf{return} V$ , then also  $P' p[P'] \rightsquigarrow^* \mathbf{return} V$ .

To show this, we note that the tree  $t$  provides an analysis of the reduction behaviour of  $P p[P]$ , and this behaviour can be seen to be mimicked by  $P' p[P']$  using the induction hypothesis together with the fact that  $P'$  has tree  $t'$  and  $p[P]$  does not denote the point  $\pi$ .

From the above claim one may now read off that if  $C P \rightsquigarrow^* \mathbf{return} c$  then also  $C P' \rightsquigarrow^* \mathbf{return} c$ . This gives the desired contradiction, as we have already noted that these values must be different.  $\square$

**COROLLARY 5.12.** *Suppose  $C$  and  $P$  are as in the preceding Lemma. For any semantic  $n$ -point  $\pi$ , the reduction sequence for  $C P$  contains at least  $n$  occurrences of terms  $\mathcal{F}[p i]$ , where  $\mathcal{F}[-]$  is an evaluation context,  $p$  is an  $n$ -point denoting  $\pi$ , and  $i$  is a natural number value.*

PROOF. Let  $\pi$  be any semantic  $n$ -point. By the previous lemma, the reduction sequence for  $C P$  contains some term  $\mathcal{E}[P p]$  where  $p$  is an  $n$ -point denoting  $\pi$ ; and the  $n$ -standardness of  $P$  tells us that the reduction sequence for  $P p$  contains  $n$  occurrences of terms  $\mathcal{G}[p i]$  where  $i$  is a natural number value and  $\mathcal{G}$  is an evaluation context. Hence the reduction sequence for  $C P$  contains  $n$  occurrences of terms  $\mathcal{F}[p i] \equiv \mathcal{E}[\mathcal{G}[p i]]$ .  $\square$

The desired lower bound now follows. Since our  $n$ -points  $p$  are assumed to be values, it is clearly impossible for the same term to be of the form  $\mathcal{E}[p i]$  and  $\mathcal{E}'[p' i']$  for two distinct  $n$ -points  $p, p'$  and evaluation contexts  $\mathcal{E}, \mathcal{E}'$ . It is therefore immediate from our corollary that the reduction sequence for  $C P$  consists of at least  $n2^n$  distinct terms, i.e. the reduction has length  $\geq n2^n$ .

**THEOREM 5.13.** *If  $C$  is an  $n$ -count program and  $P$  is any  $n$ -standard predicate, then the evaluation of  $C P$  must take time  $\Omega(n2^n)$ .*  $\square$

As we shall see, the above argument goes through with just minor adjustments for an extension of  $\lambda_b$  with exceptions, and also for a language containing the memoisation primitive required for BergerCount. For a stateful language, however, some further ingredients are required: we will return to this in Section 6.

## 6 EXTENSIONS AND VARIATIONS

Our complexity result is robust in that continues to hold in more general settings. We outline here how it generalises beyond  $n$ -standard predicates and to richer base languages.

## 6.1 Beyond $n$ -Standard Predicates

The  $n$ -standard restriction on predicates serves to make the efficiency phenomenon stand out as clearly as possible. However, we can relax the restriction by tweaking effcount to handle repeated queries and missing queries. The trade off is that the analysis of effcount becomes more involved. The key to relaxing the  $n$ -standard restriction is the use of state to keep track of which queries have been computed. We can give stateful implementations of effcount without changing its type signature by using *parameter-passing* [Kammar et al. 2013; Pretnar 2015] to internalise state within a handler. Parameter-passing abstracts every handler clause such that the current state is supplied before evaluation of a clause continues and the state is threaded through resumptions: a resumption becomes a two-argument curried function  $r : B \rightarrow S \rightarrow D$ , where the first argument of type  $B$  is the return type of the operation and the second argument is the updated state of type  $S$ .

*Repeated queries.* We can generalise effcount to handle repeated queries by memoising previous answers. First, we generalise the type of Branch such that it carries an index of a query.

$$\text{Branch}_n : \text{Nat}_n \rightarrow \text{Bool}$$

We assume a family of natural number to boolean maps,  $\text{Map}_n$  with the following interface.

$$\begin{aligned} \text{empty}_n &: \text{Map}_n \\ \text{add}_n &: \langle \text{Nat}_n, \text{Bool} \rangle \rightarrow \text{Map}_n \rightarrow \text{Map}_n \\ \text{lookup}_n &: \text{Nat}_n \rightarrow \text{Map}_n \rightarrow 1 + \text{Bool} \end{aligned}$$

Invoking the lookup function  $\text{lookup } i \text{ map}$  returns **inl**  $\langle \rangle$  if  $i$  is not present in  $\text{map}$ , and **inr**  $\text{ans}$  if  $i$  is present, where  $\text{ans} : \text{Bool}$  is the value associated with  $i$ . We can realise suitable maps in  $\lambda_b$  such that the time complexity of  $\text{add}_n$  and  $\text{lookup}_n$  is  $\mathcal{O}(\log n)$  [Okasaki 1999].

We can now use parameter-passing to support repeated queries as follows.

$$\begin{aligned} \text{effcount}'_n &: ((\text{Nat}_n \rightarrow \text{Bool}) \rightarrow \text{Bool}) \rightarrow \text{Nat} \\ \text{effcount}'_n P &:= \mathbf{let } h \leftarrow \mathbf{handle } P(\lambda i.\mathbf{do } \text{Branch}_n i) \mathbf{with} \\ \mathbf{val } b &\quad \mapsto \lambda s.\mathbf{if } b \mathbf{then return } 1 \mathbf{else return } 0 \\ \text{Branch}_n i r &\mapsto \lambda s.\mathbf{case } \text{lookup}_n i s \{ \\ &\quad \mathbf{inl } \langle \rangle \mapsto \mathbf{let } x_{\text{true}} \leftarrow r \text{ true } (\text{add}_n \langle i, \text{true} \rangle s) \mathbf{in} \\ &\quad \quad \mathbf{let } x_{\text{false}} \leftarrow r \text{ false } (\text{add}_n \langle i, \text{false} \rangle s) \mathbf{in} \\ &\quad \quad \mathbf{return } (x_{\text{true}} + x_{\text{false}}); \\ &\quad \mathbf{inr } b \mapsto r b s \} \\ \mathbf{in } h \text{ empty}_n \end{aligned}$$

The state parameter  $s$  memoises query results, thus avoiding double-counting and enabling  $\text{effcount}'_n$  to work correctly for predicates performing the same query multiple times.

*Missing queries.* Similarly, we can use parameter-passing to support missing queries.

$$\begin{aligned} \text{effcount}''_n &: ((\text{Nat}_n \rightarrow \text{Bool}) \rightarrow \text{Bool}) \rightarrow \text{Nat} \\ \text{effcount}''_n P &:= \mathbf{let } h \leftarrow \mathbf{handle } P(\lambda i.\mathbf{do } \text{Branch } \langle \rangle) \mathbf{with} \\ \mathbf{val } b &\quad \mapsto \lambda d.\mathbf{let } \text{result} \leftarrow \mathbf{if } b \mathbf{return } 1 \mathbf{else return } 0 \mathbf{in} \\ &\quad \quad \mathbf{return } \text{result} \times 2^{n-d} \\ \text{Branch } \langle \rangle r &\mapsto \lambda d.\mathbf{let } x_{\text{true}} \leftarrow r \text{ true } (d + 1) \mathbf{in} \\ &\quad \quad \mathbf{let } x_{\text{false}} \leftarrow r \text{ false } (d + 1) \mathbf{in} \\ &\quad \quad \mathbf{return } (x_{\text{true}} + x_{\text{false}}) \\ \mathbf{in } h 0 \end{aligned}$$

The parameter  $d$  keeps track of the current depth and the returned result is scaled up by  $2^{n-d}$  accounting for the unexplored part of the current subtree. This enables  $\text{effcount}''_n$  to operate correctly on predicates that inspect  $n$  points at most once. We leave it as an exercise for the reader

981 to combine  $\text{effcount}'_n$  and  $\text{effcount}''_n$  in order to obtain a generic count function that handles both  
 982 repeated queries and missing queries.  
 983

## 984 6.2 Extending $\lambda_b$ with State

985 Mutable state is a staple ingredient of many practical programming languages. We now outline  
 986 how our main lower bound result can be extended to a language with state. We will not give full  
 987 details, but merely point out the respects in which our previous treatment needs to be modified.

988 We have in mind an extension of  $\lambda_b$  with ML-style reference cells: we extend our grammar  
 989 for types with the new type for references  $A ::= \text{Ref } A$ , and that for computation terms with the  
 990 new forms for creating references (**letref**  $x = V$  **in**  $N$ ), dereferencing ( $!x$ ), and destructive update  
 991 ( $x := V$ ), with the familiar typing rules. We also add a new kind of value, namely *locations*  $l^A$ , of  
 992 type  $\text{Ref } A$ . We adopt a simple Scott-Strachey model of store [Scott and Strachey 1971]: a location  
 993 will be simply a natural number decorated with a type, and the execution of a stateful program  
 994 will allocate locations in the order  $0, 1, 2, \dots$ , assigning types to them as it does so. A *store*  $s$  will  
 995 be simply a type-respecting mapping from some set of locations  $\{0, \dots, l-1\}$  to values. For the  
 996 purposes of small-step operational semantics, a *configuration* will be a triple  $(M, l, s)$ , where  $M$  is  
 997 a computation,  $l$  is a 'location counter', and  $s$  is a store with domain  $\{0, \dots, l-1\}$ . A reduction  
 998 relation  $\rightsquigarrow$  on configurations is defined in a familiar way (again we omit the details). We shall refer  
 999 to the resulting stateful language as  $\lambda_s$ .

1000 Certain aspects of our setup require care in the presence of state. For instance, there is in general  
 1001 no unique way to assign an (untimed) decision tree to a closed value  $P : \text{Predicate}_n$ , since the  
 1002 behaviour of  $P$  on a value  $p : \text{Point}_n$  may depend both on the initial state when  $P$  is invoked, and  
 1003 on the ways in which the associated computations  $p V \rightsquigarrow^* \text{return } W$  modify the state. In this  
 1004 situation, there is not even a clear specification for what an  $n$ -count program ought to do.

1005 The simplest way to circumvent this difficulty is to define a predicate to be a closed value  
 1006  $P : \text{Predicate}_n$  within the *sublanguage*  $\lambda_b$ . For such predicates, the notions of decision tree, counting  
 1007 function and  $n$ -standardness are unproblematic. Our result will establish a runtime lower bound of  
 1008  $\Omega(n2^n)$  for count programs  $C \in \lambda_s$  applied only to predicates  $P$  of this kind.

1009 On the other hand, since  $C$  itself may be stateful, we cannot exclude the possibility that  $C P$  will  
 1010 apply  $P$  to terms  $p$  that are themselves stateful. Such a term  $p$  will no longer unambiguously denote  
 1011 some semantic point  $\pi$ , and this means the proof of Section 5.6 will not go through as it stands.

1012 To adapt our proof to the setting of  $\lambda_s$ , a little more machinery will be helpful. If  $C$  is an  $n$ -count  
 1013 program and  $P$  an  $n$ -standard predicate, we expect that the evaluation of  $C P$  will at various points  
 1014 feature terms  $\mathcal{E}[P p]$  which are then reduced in subsequent steps to some  $\mathcal{E}[\text{return } W]$ , via a  
 1015 reduction sequence which, modulo  $\mathcal{E}[-]$ , has the following form:

$$1016 \quad P p \rightsquigarrow^* \mathcal{E}_0[p i_0] \rightsquigarrow^* \mathcal{E}_0[\text{return } b_0] \rightsquigarrow^* \dots \rightsquigarrow^* \mathcal{E}_{n-1}[p i_{n-1}] \rightsquigarrow^* \mathcal{E}_{n-1}[\text{return } b_{n-1}] \rightsquigarrow^* \text{return } W$$

1017  
 1018  
 1019 (For notational clarity, we suppress mention of the location and store components here.) Informally,  
 1020 one can think of this as a dialogue in which control passes back and forth between  $P$  and  $p$ . We  
 1021 shall refer to the portions  $\mathcal{E}_j[p i_j] \rightsquigarrow^* \mathcal{E}_j[\text{return } b_j]$  of the above reduction as *p-sections*, and to the  
 1022 remaining portions (including the first and the last) as *P-sections*. We refer to the totality of these  
 1023 *P-sections* and *p-sections* as the *thread* arising from the given occurrence of the application  $P p$ .  
 1024 An important point to note is that since  $p$  may contain other occurrences of  $P$ , it is quite possible  
 1025 for the *p-sections* above to contain further threads corresponding to other applications  $P p'$ .

1026 Since  $P$  is  $n$ -standard, we know that each thread will consist of  $n + 1$  *P-sections* separated by  $n$   
 1027 *p-sections*. Indeed, it is clear that this computation traces the path  $b_0 \dots b_{n-1}$  through the decision  
 1028 tree for  $P$ , with  $i_0, \dots, i_{n-1}$  the corresponding internal node labels. We may now construe  $b_0 \dots b_{n-1}$   
 1029

1030 as a semantic point  $\pi : \mathbb{N}_n \rightarrow \mathbb{B}$ , and call it the semantic point associated with (the thread arising  
1031 from) the application occurrence  $P p$ .

1032 The following lemma now serves as a surrogate for Lemma 5.11:

1033  
1034 **LEMMA 6.1.** *Let  $P$  be an  $n$ -standard predicate. For any semantic point  $\pi : \mathbb{N}_n \rightarrow \mathbb{B}$ , the evaluation  
1035 of  $C P$  involves an application occurrence  $P p$  associated with  $\pi$ .*

1036  
1037 The proof of this lemma is not too different from that of Lemma 5.11: if  $\pi$  were a point with no  
1038 associated thread, there would be an unvisited leaf in the decision tree, and we could manufacture  
1039 an  $n$ -standard predicate  $P'$  whose tree differed from that of  $P$  only at this leaf. We can then show,  
1040 by induction on length of reductions, that any portion of the evaluation of  $C P$  can be suitably  
1041 mimicked with  $P$  replaced by  $P'$ . Naturally, this idea now needs to be formulated at the level of  
1042 configurations rather than plain terms: in the course of reducing  $(C P, 0, [])$ , we may encounter  
1043 configurations  $(M, l, s)$  in which residual occurrences of  $P$  have found their way into  $s$  as well as  $M$ ,  
1044 so in order to replace  $P$  by  $P'$  we must abstract on all these occurrences via an evident notion of  
1045 *configuration context*. With this adjustment, however, the argument of Lemma 5.11 goes through.

1046 Since each thread involves at least the  $n$  terms  $\mathcal{E}_j[p i_j]$ , our proof of the  $\Omega(n2^n)$  bound is complete  
1047 provided we can show that no two threads overlap: more precisely, none of the above terms  $\mathcal{E}_j[p i_j]$   
1048 can belong to the  $P$ -section of more than one thread. The difficulty here is that because syntactic  
1049 points no longer have unambiguous denotations, the relevant  $\pi$  can no longer be simply read off  
1050 from  $p$ : indeed, it is entirely possible that our computation may involve two instances of the same  
1051 application  $P p$  giving rise to entirely different threads owing to the presence of state. Fortunately,  
1052 however, we may reason as follows.

1053 Let us suppose that  $P p$  and  $P p'$  are any two application occurrences arising in the evaluation of  
1054  $C P$ , with  $P p$  appearing before  $P p'$ , and suppose these respectively give rise to threads  $T, T'$ . We  
1055 wish to show that the  $P$ -sections of  $T$  do not overlap with those of  $T'$ . There are three cases:

- 1056 • If  $T'$  does not start until after  $T$  has finished, then of course  $T, T'$  are disjoint.
- 1057 • If  $T'$  starts within some  $p$ -section  $\mathcal{E}_j[p i_j] \rightsquigarrow^* \mathcal{E}_j[\mathbf{return} b_j]$  of  $T$ , then it is not hard to see  
1058 that  $T'$  must also end within this same  $p$ -section, as the evaluation of  $P p'$  forms part of the  
1059 evaluation of  $p i_j$ .
- 1060 • It is not possible for  $T'$  to start within a  $P$ -section of  $T$ . This follows from the fact that a  
1061 ‘residual occurrence’ of  $P$  (that is, one arising as a residual of the  $P$  in  $C P$ ) cannot itself  
1062 contain other residual occurrences of  $P$ ; thus, for any term arising from the reduction of  $P p$   
1063 (discounting  $P p$  itself), every residual occurrence of  $P$  occurs within some  $p$ .

1064 Arguing along such lines, one can show that any two threads are indeed ‘disjoint’, in such a way  
1065 that there must be at least  $n2^n$  steps in the overall reduction sequence.

### 1067 6.3 Berger Count

1068  
1069 We now briefly outline the BergerCount program alluded to in Section 5.6, in order to fill out our  
1070 overall picture of the relationship between language expressivity and potential program efficiency.

1071 Berger’s original program [Berger 1990] introduced a remarkable search operator for predicates  
1072 on *infinite* streams of booleans, and has played an important role in higher-order computability  
1073 theory [Longley and Normann 2015]. What we wish to highlight here is that if one applies the algo-  
1074 rithm to predicates on *finite* boolean vectors, the resulting program, though no longer interesting  
1075 from a computability perspective, still holds some interest from a complexity standpoint: indeed, it  
1076 yields what seems to be the best available implementation of generic counting within a PCF-style  
1077 ‘functional’ language (provided one accepts the use of a primitive for call-by-need evaluation).

We give the gist of an adaptation of Berger's search algorithm on finite spaces.

```

1079
1080     bestshotn : Predicaten → Pointn
1081     bestshotn P := bestshot'n P []
1082
1083     bestshot'n : Predicaten → List Bool → Pointn
1084     bestshot'n P start := let f ← memoise (λ⟨⟩.bestshot''n P start) in
1085                          return (λi.if i < |start| then start.i else (f ⟨⟩).i)
1086
1087     bestshot''n : Predicaten → List Bool → List Bool
1088     bestshot''n P start := if |start| = n then return start
1089                          else let f ← bestshot'n P (append start [true]) in
1090                             if P f then return [f 0, ..., f (n - 1)]
1091                             else bestshot''n P (append start [false])

```

The function  $\text{bestshot}_n$  will return a point satisfying the given predicate  $P$  if there is one, or the dummy point  $\lambda i.\text{false}$  if there is none. This is implemented by means of two mutually recursive auxiliary functions whose workings are admittedly hard to elucidate in a few words. The function  $\text{bestshot}'_n$  is a generalisation of  $\text{bestshot}_n$  that makes a best shot at finding a point  $p$  satisfying  $P$  and matching some specified list  $\text{start}$  in some initial segment of its components  $[p\ 0, \dots, p\ (i-1)]$ . This works 'lazily', drawing its values from  $\text{start}$  wherever possible, and performing an actual search only when required. This actual search is undertaken by  $\text{bestshot}''_n$ , which proceeds by first searching for a solution that extends the specified list with true; but if no such solution is forthcoming, it settles for false as the next component of the point being constructed. The whole procedure relies on a subtle combination of laziness, recursion and implicit nesting of calls to  $P$  which means that the search is self-pruning in regions of the binary tree where  $P$  only demands some initial segment  $p\ 0, \dots, p\ (i-1)$  of its argument  $p$ .

The above program makes use of an operation

$$\text{memoise} : (1 \rightarrow \text{List Bool}) \rightarrow (1 \rightarrow \text{List Bool})$$

which transforms a given thunk into an equivalent 'memoised' version, i.e. one that caches its value after its first invocation and immediately returns this value on all subsequent invocations. Such an operation may readily be implemented in  $\lambda_s$ , or alternatively may simply be added as a primitive in its own right (we omit the details). The latter has the advantage that it preserves the purely 'functional' character of the language, in the sense that every program is observationally equivalent to a  $\lambda_b$  program, namely the one obtained by replacing  $\text{memoise}$  by the identity.

We now show how the above idea may be exploited to yield a generic count program (this part of our work appears to be new).

```

1113     BergerCountn : Predicaten → Nat
1114     BergerCountn P := count'n P [] 0
1115
1116     count'n : Predicaten → List Bool → Nat → Nat
1117     count'n P start acc := if |start| = n then acc + (if P(λi.start.i) then return 1 else return 0)
1118                          else let f ← bestshot'n P start |start| in
1119                             if P f then count''n start [f 0, ..., f (n - 1)] acc else return acc
1120
1121     count''n : Predicaten → List Bool → List Bool → Nat → Nat
1122     count''n P start leftmost acc := if |start| = n then acc + 1
1123                          else let b ← leftmost.start | in
1124                             let acc' ← count''n (append start [b]) leftmost acc in
1125                             if b then count''n (append start [false]) acc' else return acc'

```

Again,  $\text{BergerCount}_n$  is implemented by means of two mutually recursive auxiliary functions. The function  $\text{count}'_n$  counts the solutions to  $P$  that start with the specified list of booleans, adding their number to a previously accumulated total given by  $\text{acc}$ . The function  $\text{count}''_n$  does the same thing,

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. Runtimes Relative to the Effectful Implementation

but exploiting the knowledge that a best shot at the ‘leftmost’ solution to  $P$  within this subtree has already been computed. (We are visualising  $n$ -points as forming a binary tree with true to the left of false at each fork.) Thus,  $\text{count}_n''$  will not re-examine the portion of the subtree to the left of this candidate solution, but rather will start at this solution and work rightward.

This gives rise to an  $n$ -count program that can work efficiently on predicates that tend to ‘fail fast’: more specifically, predicates  $P$  that inspect the components of their argument  $p$  in order  $p_0$ ,  $p_1$ ,  $p_2$ ,  $\dots$ , and which are frequently able to return false after inspecting just a small number of these components. Generalising our program from binary to  $k$ -ary branching trees, we see that the  $n$ -queens problem provides a typical example: most points in the space can be seen *not* to be solutions by inspecting just the first few components. Our experimental results in Section 7 attest to the viability of this approach and its overwhelming superiority over the naive functional method.

By contrast, the above program is *not* able to take advantage of parts of the tree where our predicate ‘succeeds fast’, i.e. returns true after seeing only a few components. Unlike the effectful count program of Section 5.2, which may sometimes add  $2^{n-d}$  to the count in a single step, the Berger approach can only count solutions one at a time. Thus, any evaluation of  $\text{count}_n P$  that returns a natural number  $c$  must take time  $\Omega(c)$ . These observations informally indicate the likely extent of the efficiency gap between effectful and purely functional computation when it comes to non- $n$ -standard predicates.

## 7 EXPERIMENTS

The theoretical efficiency gap between realisations of  $\lambda_b$  and  $\lambda_h$  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  $\lambda_b$  extended with mutable state:

- Naïve: a simple, and rather naïve, functional implementation;
- Pruned: a generic search procedure with space pruning based on Longley’s technique [Longley 1999] (uses local state);
- Berger: a lazy pure functional generic search procedure based on Berger’s algorithm.

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  $\infty$ . 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



1177 implementation uses an encoding of delimited control akin to effect handlers based on top of  
 1178 SML/NJ's call/cc.

1179

1180 *Queens.* We phrase the  $n$ -queens problem as a generic search problem. As a control we include a  
 1181 bespoke implementation hand-optimised for the problem. We perform two experiments: finding  
 1182 the first solution for  $n \in \{20, 24, 28\}$  and enumerating all solutions for  $n \in \{8, 10, 12\}$ . The speedup  
 1183 over the naïve implementation is dramatic, but less so over the Berger procedure. The pruned  
 1184 procedure is more competitive, but still slower than the baseline. Unsurprisingly, the baseline is  
 1185 slower than the bespoke implementation.

1186

1187 *Exact Real Integration.* The integration benchmarks are adapted from Simpson [1998]. We inte-  
 1188 grate three different functions with varying precision in the interval  $[0, 1]$ . For the identity function  
 1189 (Id) at precision 20 the speedup relative to Berger is  $5.59\times$ . For the squaring function the speedups  
 1190 are larger at higher precisions: at precision 14 the speedup is  $4.39\times$  over the pruned integrator,  
 1191 whilst it is  $5.08\times$  at precision 20. The speedups are more extreme against the naïve and Berger  
 1192 integrators. We also integrate the logistic map  $x \mapsto 1 - 2x^2$  at a fixed precision of 15. We make  
 1193 the function harder to compute by iterating it up to 5 times. Between the pruned and effectful  
 1194 integrator the speedup ratio increases as the function becomes harder to compute.

1195

1196

1197

## 8 CONCLUSIONS AND FUTURE WORK

1198 We presented a PCF-inspired language  $\lambda_b$  and its extension with effect handlers  $\lambda_h$ . We proved that  
 1199  $\lambda_h$  exhibits an asymptotically more efficient implementation of generic search than any possible  
 1200 implementation in  $\lambda_b$ . We observed its effect in practice on several benchmarks. We also proved  
 1201 that our  $\Omega(n2^n)$  lower bound applies to a language  $\lambda_s$  which extends  $\lambda_b$  with state.

1202 We have also verified that the same lower bound applies to a language  $\lambda_e$  which extends  $\lambda_b$  with  
 1203 (Benton-Kennedy style [Benton and Kennedy 2001]) *exceptions* and handlers — and even for the  
 1204 combined language  $\lambda_{se}$  with both state and exceptions. As was the case for  $\lambda_s$ , it is helpful to insist  
 1205 here that our predicates themselves are terms of  $\lambda_b$ . However, the adaptations of our proof method  
 1206 required for  $\lambda_e$  are less interesting and far-reaching than those for  $\lambda_s$  so we do not present them  
 1207 here. We also remark informally that  $\lambda_{se}$  seems to bring us close to the expressive power of real  
 1208 languages such as Standard ML, Java, and Python, strongly suggesting that the speedup we have  
 1209 discussed is unattainable in these language.

1210 The result extends to other control operators by appeal to existing results on interdefinability of  
 1211 handlers and other control operators [Forster et al. 2019; Piróg et al. 2019]. The result no longer  
 1212 applies directly if we add an effect type system to  $\lambda_h$ , as the implementation of the counting program  
 1213 would require a change of type for predicates to reflect the ability to perform effectful operations.  
 1214 In future we plan to investigate how to account for effect type systems.

1215 One might object that the efficiency gap we have analysed is of merely theoretical interest,  
 1216 since an  $\Omega(2^n)$  runtime is already ‘infeasible’. What we claim, however, is that what we have  
 1217 presented is an example of a much more pervasive phenomenon, and our generic counting example  
 1218 serves merely as a convenient way to bring this phenomenon into sharp formal focus. Suppose, for  
 1219 example, that our programming task was not to count all solutions to  $P$ , but to find just one of them.  
 1220 It is informally clear that for many kinds of predicates this would in practice be a feasible task, and  
 1221 also that we could still gain our factor  $n$  speedup here by working in a language with first-class  
 1222 control. However, such an observation appears less amenable to a clean mathematical formulation,  
 1223 as the runtimes in question are highly sensitive to both the particular choice of predicate and the  
 1224 search order employed.

1225

## 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

- Andrej Bauer. 2018. What is algebraic about algebraic effects and handlers? *CoRR* abs/1807.05923 (2018).
- Andrej Bauer and Matija Pretnar. 2015. Programming with algebraic effects and handlers. *J. Log. Algebr. Meth. Program.* 84, 1 (2015), 108–123.
- Jordan Bell and Brett Stevens. 2009. A survey of known results and research areas for n-queens. *Discret. Math.* 309, 1 (2009), 1–31.
- Nick Benton and Andrew Kennedy. 2001. Exceptional Syntax Journal of Functional Programming. *J. Funct. Program.* 11, 4 (2001), 395–410.
- Ulrich Berger. 1990. *Totale Objekte und Mengen in der Bereichstheorie*. Ph.D. Dissertation. Ludwig Maximilians-Universität, Munich.
- Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Abstracting algebraic effects. *PACMPL* 3, POPL (2019), 6:1–6:28.
- Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2020. Binders by day, labels by night: effect instances via lexically scoped handlers. *PACMPL* 4, POPL (2020), 48:1–48:29.
- 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.
- Richard S. Bird. 2006. Functional Pearl: A program to solve Sudoku. *J. Funct. Program.* 16, 6 (2006), 671–679.
- Robert Cartwright and Matthias Felleisen. 1992. Observable Sequentiality and Full Abstraction. In *POPL*. ACM Press, 328–342.
- Lukas Convent, Sam Lindley, Conor McBride, and Craig McLaughlin. 2020. Doo bee doo bee doo. *J. Funct. Program.* 30 (2020). To appear.
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms, Third Edition* (3rd ed.). MIT Press.
- 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](http://homepages.inf.ed.ac.uk/jrl/Research/Robbie_Daniels_MSc_dissertation.pdf)
- Olivier Danvy and Andrzej Filinski. 1990. Abstracting Control. In *LISP and Functional Programming*. ACM, 151–160.
- Stephen Dolan, Leo White, KC Sivaramakrishnan, Jeremy Yallop, and Anil Madhavapeddy. 2015. Effective Concurrency through Algebraic Effects. OCamL Workshop. (2015).
- Martín Hötzel Escardó. 2007. Infinite sets that admit fast exhaustive search. In *LICS*. IEEE Computer Society, 443–452.
- 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. AAI8727494.
- Matthias Felleisen. 1988. The Theory and Practice of First-Class Prompts. In *POPL*. ACM Press, 180–190.
- Matthias Felleisen. 1991. On the expressive power of programming languages. *Sci. Comput. Prog.* 17, 1–3 (1991), 35–75.
- Matthias Felleisen and Daniel P. Friedman. 1987. Control Operators, the SECD-machine, and the  $\lambda$ -Calculus. In *The Proceedings of the Conference on Formal Description of Programming Concepts III, Eberup, Denmark*. Elsevier, 193–217.
- Andrzej Filinski. 1994. Representing Monads. In *POPL*. ACM Press, 446–457.
- Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. 1993. The Essence of Compiling with Continuations. In *PLDI*. ACM, 237–247.
- Yannick Forster, Ohad Kammar, Sam Lindley, and Matija Pretnar. 2019. On the expressive power of user-defined effects: Effect handlers, monadic reflection, delimited control. *J. Funct. Program.* 29 (2019), e15.
- Daniel Hillerström and Sam Lindley. 2016. Liberating effects with rows and handlers. In *TyDe@ICFP*. ACM, 15–27.
- Daniel Hillerström and Sam Lindley. 2018. Shallow Effect Handlers. In *APLAS (Lecture Notes in Computer Science)*, Vol. 11275. Springer, 415–435.
- Daniel Hillerström, Sam Lindley, and Robert Atkey. 2020. Effect Handlers via Generalised Continuations. *J. Funct. Program.* 30 (2020). To appear.
- 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.
- Neil Jones. 2001. The expressive power of higher-order types, or, life without CONS. *J. Funct. Progr.* 11 (2001), 5–94.
- Ohad Kammar, Sam Lindley, and Nicolas Oury. 2013. Handlers in action. In *ICFP*. ACM, 145–158.

- 1275 Oleg Kiselyov, Amr Sabry, and Cameron Swords. 2013. Extensible effects: an alternative to monad transformers. In *Haskell*.  
1276 ACM, 59–70.
- 1277 Oleg Kiselyov, Chung-chieh Shan, Daniel P. Friedman, and Amr Sabry. 2005. Backtracking, Interleaving, and Terminating  
1278 Monad Transformers: (Functional Pearl). (2005), 192–203.
- 1279 Donald Knuth. 1997. *The Art of Computer Programming, Volume 1: Fundamental Algorithms (third edition)*. Addison-Wesley.
- 1280 Daan Leijen. 2017. Type directed compilation of row-typed algebraic effects. In *POPL*. ACM, 486–499.
- 1281 Paul Blain Levy, John Power, and Hayo Thielecke. 2003. Modelling environments in call-by-value programming languages.  
1282 *Inf. Comput.* 185, 2 (2003), 182–210.
- 1283 Sam Lindley. 2014. Algebraic effects and effect handlers for idioms and arrows. In *WGP@ICFP*. ACM, 47–58.
- 1284 Sam Lindley, Conor McBride, and Craig McLaughlin. 2017. Do be do be do. In *POPL*. ACM, 500–514.
- 1285 John Longley. 1999. When is a functional program not a functional program?. In *ICFP*. ACM, 1–7.
- 1286 John Longley. 2018. The recursion hierarchy for PCF is strict. *Logical Methods in Comput. Sci.* 14, 3:8 (2018), 1–51.
- 1287 John Longley. 2019. Bar recursion is not computable via iteration. *Computability* 8, 2 (2019), 119–153.
- 1288 John Longley and Dag Normann. 2015. *Higher-Order Computability*. Springer.
- 1289 Robin Milner. 1977. Fully Abstract Models of Typed  $\lambda$ -Calculi. *Theor. Comput. Sci.* 4, 1 (1977), 1–22.
- 1290 MLton. 2020. MLton website. (2020). <http://www.mlton.org>
- 1291 Chris Okasaki. 1999. *Purely functional data structures*. Cambridge University Press.
- 1292 Nicholas Pippenger. 1996. Pure versus impure Lisp. In *POPL*. ACM, 104–109.
- 1293 Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Typed Equivalence of Effect Handlers and Delimited Control. In  
1294 *FSCD (LIPIcs)*, Vol. 131. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 30:1–30:16.
- 1295 Gordon Plotkin. 1977. LCF considered as a programming language. *Theor. Comput. Sci.* 5, 3 (1977), 223–255.
- 1296 Gordon D. Plotkin and John Power. 2001. Adequacy for Algebraic Effects. In *FoSSaCS (Lecture Notes in Computer Science)*,  
1297 Vol. 2030. Springer, 1–24.
- 1298 Gordon D. Plotkin and Matija Pretnar. 2013. Handling Algebraic Effects. *Logical Methods in Computer Science* 9, 4 (2013).
- 1299 Matija Pretnar. 2015. An Introduction to Algebraic Effects and Handlers. *Electr. Notes Theor. Comput. Sci.* 319 (2015), 19–35.  
1300 Invited tutorial paper.
- 1301 Dana Scott and Christopher Strachey. 1971. *Proceedings of the Symposium on Computers and Automata* 21 (1971).
- 1302 Alex K. Simpson. 1998. Lazy Functional Algorithms for Exact Real Functionals. In *MFCs (Lecture Notes in Computer Science)*,  
1303 Vol. 1450. Springer, 456–464.
- 1304 Michael Sperber, Kent R. Dybvig, Matthew Flatt, Anton van Stratten, Robby Bruce Findler, and Jacob Matthews. 2009.  
1305 Revised<sup>6</sup> Report on the Algorithmic Language Scheme. *J. Funct. Progr.* 19, S1 (2009), 1–301.

1301  
1302  
1303  
1304  
1305  
1306  
1307  
1308  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1320  
1321  
1322  
1323

Configurations	Pure continuations
$\langle\langle M \mid \gamma \mid \sigma \rangle\rangle = \langle\langle \sigma \rangle\rangle(\langle\langle M \rangle\rangle\gamma)$	$\langle\langle [] \rangle\rangle M = M$
	$\langle\langle (\gamma, x, N) :: \sigma \rangle\rangle M = \langle\langle \sigma \rangle\rangle(\mathbf{let} \ x \leftarrow M \ \mathbf{in} \ \langle\langle N \rangle\rangle(\gamma \setminus \{x\}))$
Computation terms	
	$\langle\langle (V \ W) \rangle\rangle\gamma = \langle\langle V \rangle\rangle\gamma \ \langle\langle W \rangle\rangle\gamma$
	$\langle\langle (\mathbf{let} \ \langle x; y \rangle = V \ \mathbf{in} \ N) \rangle\rangle\gamma = \mathbf{let} \ \langle x; y \rangle = \langle\langle V \rangle\rangle\gamma \ \mathbf{in} \ \langle\langle N \rangle\rangle(\gamma \setminus \{x, y\})$
	$\langle\langle (\mathbf{case} \ V \ \{\mathbf{inl} \ x \mapsto M; \mathbf{inr} \ y \mapsto N\}) \rangle\rangle\gamma = \mathbf{case} \ \langle\langle V \rangle\rangle\gamma \ \{\mathbf{inl} \ x \mapsto \langle\langle M \rangle\rangle(\gamma \setminus \{x\});$
	$\mathbf{inr} \ y \mapsto \langle\langle N \rangle\rangle(\gamma \setminus \{y\})\}$
	$\langle\langle (\mathbf{return} \ V) \rangle\rangle\gamma = \mathbf{return} \ \langle\langle V \rangle\rangle\gamma$
	$\langle\langle (\mathbf{let} \ x \leftarrow M \ \mathbf{in} \ N) \rangle\rangle\gamma = \mathbf{let} \ x \leftarrow \langle\langle M \rangle\rangle\gamma \ \mathbf{in} \ \langle\langle N \rangle\rangle(\gamma \setminus \{x\})$
Value terms and values	
$\langle\langle x \rangle\rangle\gamma = \langle\langle v \rangle\rangle$ , if $\gamma(x) = v$	$\langle\langle n \rangle\rangle = n$
$\langle\langle x \rangle\rangle\gamma = x$ , if $x \notin \text{dom}(\gamma)$	$\langle\langle (\gamma, \lambda x^A. M) \rangle\rangle = \lambda x^A. \langle\langle M \rangle\rangle(\gamma \setminus \{x\})$
$\langle\langle n \rangle\rangle\gamma = n$	$\langle\langle (\gamma, \mathbf{rec} \ f \ x^A. M) \rangle\rangle = \mathbf{rec} \ f \ x^A. \langle\langle M \rangle\rangle(\gamma \setminus \{f, x\})$
$\langle\langle \lambda x^A. M \rangle\rangle\gamma = \lambda x^A. \langle\langle M \rangle\rangle(\gamma \setminus \{x\})$	$\langle\langle \langle \rangle \rangle = \langle \rangle$
$\langle\langle (\mathbf{rec} \ f \ x^A. M) \rangle\rangle\gamma = \mathbf{rec} \ f \ x^A. \langle\langle M \rangle\rangle(\gamma \setminus \{f, x\})$	$\langle\langle \langle v; w \rangle \rangle = \langle \langle v \rangle \rangle; \langle \langle w \rangle \rangle$
$\langle\langle \langle \rangle \rangle\rangle\gamma = \langle \rangle$	$\langle\langle (\mathbf{inl} \ v)^B \rangle\rangle = \langle \mathbf{inl} \ \langle \langle v \rangle \rangle \rangle^B$
$\langle\langle \langle V; W \rangle \rangle\rangle\gamma = \langle \langle \langle V \rangle \rangle \rangle\gamma; \langle \langle \langle W \rangle \rangle \rangle\gamma$	$\langle\langle (\mathbf{inr} \ w)^A \rangle\rangle = \langle \mathbf{inr} \ \langle \langle w \rangle \rangle \rangle^A$
$\langle\langle (\mathbf{inl} \ V)^B \rangle\rangle\gamma = \langle \mathbf{inl} \ \langle \langle V \rangle \rangle \rangle^B$	$\langle\langle \langle \sigma^A \rangle \rangle = \lambda x^A. \langle \langle \sigma \rangle \rangle(\mathbf{return} \ x)$
$\langle\langle (\mathbf{inr} \ W)^A \rangle\rangle\gamma = \langle \mathbf{inr} \ \langle \langle W \rangle \rangle \rangle^A$	

Fig. 8. Mapping from Base Machine Configurations to Terms

## A CORRECTNESS OF THE BASE MACHINE

We now show that the base 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 8 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,  $\gamma$ , which we define now.

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

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  $\beta$ -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 (M-LET) and  $\simeq_a$  for the symmetric closure of  $\longrightarrow_a^*$ . We write  $\longrightarrow_\beta$  for  $\beta$ -steps (all other rules) 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  $\beta$ -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  $M \rightsquigarrow N$ . □

1373	Configurations	Continuations
1374	$\langle\langle M \mid \gamma \mid \kappa \rangle\rangle = \langle\langle \kappa \rangle\rangle(\langle\langle M \rangle\rangle\gamma)$	$\langle\langle [] \rangle\rangle M = M$
1375		$\langle\langle (\sigma, \chi) :: \kappa \rangle\rangle M = \langle\langle \kappa \rangle\rangle(\langle\langle \chi \rangle\rangle(\langle\langle \sigma \rangle\rangle(M)))$
1376	Handler Closures and Definitions	
1377	$\langle\langle \gamma, H \rangle\rangle M = \mathbf{handle} M \mathbf{with} \langle\langle H \rangle\rangle\gamma$	$\langle\langle \{\mathbf{val} x \mapsto M\} \rangle\rangle\gamma = \{\mathbf{val} x \mapsto \langle\langle M \rangle\rangle(\gamma \setminus \{x\})\}$
1378		$\langle\langle \{\ell x r \mapsto M\} \uplus H \rangle\rangle\gamma = \{\ell x r \mapsto \langle\langle M \rangle\rangle(\gamma \setminus \{x, r\})\} \uplus \langle\langle H \rangle\rangle\gamma$
1379	Computation Terms and Machine Values	
1380	$\langle\langle \mathbf{handle} M \mathbf{with} H \rangle\rangle\gamma = \mathbf{handle} \langle\langle M \rangle\rangle\gamma \mathbf{with} \langle\langle H \rangle\rangle\gamma$	$\langle\langle \gamma, H \rangle\rangle\gamma = \lambda x^A. \langle\langle \gamma, H \rangle\rangle(\mathbf{return} x)$
1381	$\langle\langle \mathbf{do} \ell V \rangle\rangle\gamma = \mathbf{do} \ell \langle\langle V \rangle\rangle\gamma$	
1382		

Fig. 9. Mapping from Handler Machine Configurations to Terms

The correspondence here is rather strong: there is a one-to-one mapping between  $\sim$  and  $\implies / \simeq_a$ . 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 (BASE SIMULATION).** *If  $\vdash M : A$  and  $M \sim^+ N$  such that  $N$  is normal, then  $\langle M \mid \emptyset \mid [] \rangle \longrightarrow^+ C$  such that  $\langle\langle C \rangle\rangle = N$ , or  $M \not\rightsquigarrow$  then  $\langle M \mid \emptyset \mid [] \rangle \not\rightsquigarrow$ .*

**PROOF.** By repeated application of Lemma A.3. □

## B CORRECTNESS OF THE HANDLER MACHINE

The correctness result for the base machine 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. Figure 9 shows the necessary changes to the  $\langle - \rangle$  function.

**LEMMA B.1.** *Suppose  $M$  is a computation and  $C$  is configuration such that  $\langle\langle C \rangle\rangle = M$ , then if  $M \sim N$  there exists  $C'$  such that  $C \implies C'$  and  $\langle\langle C' \rangle\rangle = N$ , or if  $M \not\rightsquigarrow$  then  $C \not\rightsquigarrow$ .*

**PROOF.** By induction on the derivation of  $M \sim N$ . □

**THEOREM B.2 (HANDLER SIMULATION).** *If  $\vdash M : A$  and  $M \sim^+ N$  such that  $N$  is normal, then  $\langle M \mid \emptyset \mid \kappa_0 \rangle \longrightarrow^+ C$  such that  $\langle\langle C \rangle\rangle = N$ , or  $M \not\rightsquigarrow$  then  $\langle M \mid \emptyset \mid \kappa_0 \rangle \not\rightsquigarrow$ .*

**PROOF.** By repeated application of Lemma B.1. □

## C PROOF DETAILS FOR THE COMPLEXITY OF EFFECTFUL GENERIC COUNTING

In this appendix we give proof details and artefacts for Theorem 5.10. Throughout this section we let  $H_{\text{count}}$  denote the handler definition of count, that is

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

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

$$\mathcal{D} : \text{Conf} \rightarrow \text{DT}$$

$$\mathcal{D}(\langle \text{return true} \mid \gamma \mid [] \rangle) = (\langle \text{!true}, 0 \rangle, (\gamma, []))$$

$$\mathcal{D}(\langle \text{return false} \mid \gamma \mid [] \rangle) = (\langle \text{!false}, 0 \rangle, (\gamma, []))$$

$$\mathcal{D}(\langle p V \mid \gamma \mid \sigma \rangle) = (\langle \text{?}[V]\gamma, 0 \rangle, (\gamma, \sigma))$$

$$\mathcal{D}(\langle p V \mid \gamma \mid \sigma \rangle) (b :: bs) \simeq \mathcal{D}(\langle \text{return } b \mid \gamma \mid \sigma \rangle) bs$$

$$\mathcal{D}(\langle M \mid \gamma \mid \sigma \rangle) bs \simeq \mathcal{I}(\mathcal{D}(\langle M' \mid \gamma' \mid \sigma' \rangle) bs),$$

$$\text{if } \langle M \mid \gamma \mid \sigma \rangle \longrightarrow \langle M' \mid \gamma' \mid \sigma' \rangle$$

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  $\mathcal{D}(\langle P \mid \emptyset[p \mapsto p] \mid [] \rangle)$ .

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  $H_{\text{count}}$ .

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

*Definition C.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.

$$\text{arrive} : \mathbb{B}^* \times \text{Val} \rightarrow \text{Conf}$$

$$\text{arrive}(bs, P) := \langle V j \mid \gamma \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle, \quad \text{if } |bs| < n$$

$$\text{where } \gamma = \text{env}(t)(bs), ?j = \text{labs}(t)(bs), \text{ and } \llbracket V \rrbracket \gamma = (\text{env}^{\perp}(P), \lambda\_.\text{do Branch } \langle \rangle)$$

$$\text{arrive}(bs, P) := \langle \text{return } b \mid \gamma \mid ([], \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle, \quad \text{if } |bs| = n$$

$$\text{where } \gamma = \text{env}(t)(bs) \text{ and } !b = \text{labs}(t)(bs)$$

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

$$\text{depart} : \mathbb{B}^* \times \text{Val} \rightarrow \text{Conf}$$

$$\text{depart}(bs, P) := \langle \text{return } m \mid \gamma \mid \text{residual}(bs, P) \rangle, \quad \text{if } |bs| < n$$

$$\text{where } \gamma = \text{env}_{\text{false}}^{\uparrow}(bs, P) \text{ and } m = c(\text{true} :: bs) + c(\text{false} :: bs)$$

$$\text{depart}(bs, P) := \langle \text{return } m \mid \gamma \mid \text{residual}(bs, P) \rangle, \quad \text{if } |bs| = n$$

$$\text{where } \gamma = \text{env}^{\perp}(P) \text{ and } m = c(bs)$$

The two clauses of `depart` yield slightly different configurations. The first clause computes a configuration inside the operation clause of  $H_{\text{count}}$ . 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  $H_{\text{count}}$  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.

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

- The function  $\text{env}^\perp$  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.

$$\begin{aligned} \text{env}^\perp &: \text{Val} \rightarrow \text{Env} \\ \text{env}^\perp(P) &:= \emptyset[\text{pred} \mapsto \llbracket P \rrbracket \emptyset] \\ \text{env}_{\text{true}}^\downarrow &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Env} \\ \text{env}_{\text{true}}^\downarrow(bs, P) &:= \text{env}^\perp(V)[r \mapsto (\sigma, \chi_{\text{count}}(P))], \\ &\quad \text{where } \sigma = \text{pure}(t)(bs) \\ \text{env}_{\text{false}}^\downarrow &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Env} \\ \text{env}_{\text{false}}^\downarrow(bs, P) &:= \gamma[x \mapsto i], \\ &\quad \text{where } \gamma = \text{env}_{\text{true}}^\downarrow(bs, P) \text{ and } i = c(\text{true} :: bs) \\ \text{env}_{\text{false}}^\uparrow &: \mathbb{B}^* \times \text{Val} \rightarrow \text{Env} \\ \text{env}_{\text{false}}^\uparrow(bs, P) &:= \gamma[y \mapsto j], \\ &\quad \text{where } \gamma = \text{env}_{\text{false}}^\downarrow(bs, P) \text{ and } j = c(\text{false} :: bs) \end{aligned}$$

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  $H_{\text{count}}$ .

LEMMA C.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

(1) If  $|bs| = []$ , then

$$\begin{aligned} & \langle P \ p \ | \ \gamma \ | \ [] \rangle \\ & \longrightarrow_{\text{steps}(t)([])} \\ & \langle p \ i \ | \ \gamma' \ | \ \sigma \rangle, \end{aligned}$$

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 amount of transitions

$$\begin{aligned} & \langle P \ p \ | \ \gamma \ | \ ([], \chi_{\text{count}}(P)) :: \text{residual}(P, [])[(\lambda\_.\text{do Branch } \langle \rangle)/p] \\ & \longrightarrow_{\text{steps}(t)([])} \\ & \langle p \ i \ | \ \gamma' \ | \ (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, [])[(\lambda\_.\text{do Branch } \langle \rangle)/p] \end{aligned}$$

(2) For  $bs = b :: bs'$  we have the following two subcases

• If  $|bs| < n$ , then

$$\begin{aligned} & \langle \text{return } b \ | \ \gamma \ | \ \sigma \rangle \\ & \longrightarrow_{\text{steps}(t)(b::bs)} \\ & \langle p \ i \ | \ \gamma' \ | \ \sigma \rangle, \end{aligned}$$

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

$$\begin{aligned} & \langle \text{return } b \ | \ \gamma \ | \ (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c)[(\lambda\_.\text{do Branch } \langle \rangle)/p] \\ & \longrightarrow_{\text{steps}(t)(b::bs)} \\ & \langle p \ i \ | \ \gamma' \ | \ (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c)[(\lambda\_.\text{do Branch } \langle \rangle)/p] \end{aligned}$$

• If  $|bs| = n$ , then

$$\begin{aligned} & \langle \text{return } b \ | \ \gamma \ | \ \sigma \rangle \\ & \longrightarrow_{\text{steps}(t)(b::bs')} \\ & \langle \text{return } b' \ | \ \gamma' \ | \ [] \rangle, \end{aligned}$$

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 perform the same amount of transitions

$$\begin{aligned} & \langle \text{return } b \ | \ \gamma \ | \ (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c)[(\lambda\_.\text{do Branch } \langle \rangle)/p] \\ & \longrightarrow_{\text{steps}(t)(b::bs')} \\ & \langle \text{return } b' \ | \ \gamma' \ | \ ([], \chi_{\text{count}}(P)) :: \text{residual}(P, b :: bs, n, t, c)[(\lambda\_.\text{do Branch } \langle \rangle)/p] \end{aligned}$$

PROOF. Proof by induction on the transition relation  $\longrightarrow$ . □

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

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

The following lemma performs most of the heavy lifting for the proof of Theorem 5.10.

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

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



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

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

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.

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

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

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

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

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

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.

$$\begin{aligned}
& \text{arrive}(bs, P) \\
&= \text{(definition of arrive when } n < |bs|) \\
& \langle V j \mid \gamma \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle \\
& \text{where } ?j = \text{labs}(t)(bs), \gamma = \text{env}(t)(bs), \sigma = \text{pure}(t)(bs), \text{ and } V = (\text{env}^\perp(P), \lambda\_.\text{do Branch } \langle \rangle) \\
&\longrightarrow \text{(M-APP)} \\
& \langle \text{do Branch } \langle \rangle \mid \gamma'[_ \mapsto \llbracket j \rrbracket \gamma'] \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(bs, P) \rangle \\
& \hspace{15em} \text{where } \gamma' = \text{env}^\perp(P) \\
&\longrightarrow \text{(M-HANDLE-OP, } \chi_{\text{count}}(P)^{\text{Branch}} = \{\text{Branch } \langle \rangle r \mapsto \dots\}) \\
& \left\langle \begin{array}{l} \text{let } x_{\text{true}} \leftarrow r \text{ true in} \\ \text{let } x_{\text{false}} \leftarrow r \text{ false in } \mid \gamma[r \mapsto \llbracket (\sigma, \chi_{\text{count}}(P)) \rrbracket \gamma] \mid \text{residual}(bs, P) \end{array} \right\rangle \\
& \hspace{15em} \text{where } \gamma = \text{env}^\perp(P) \\
&= \text{(definition of } \llbracket - \rrbracket \text{ (1 value step))} \\
& \left\langle \begin{array}{l} \text{let } x_{\text{true}} \leftarrow r \text{ true in} \\ \text{let } x_{\text{false}} \leftarrow r \text{ false in } \mid \gamma' \mid \text{residual}(bs, P) \end{array} \right\rangle \\
& \hspace{15em} \text{where } \gamma' = \gamma[r \mapsto (\sigma, \chi_{\text{count}}(P))] \\
&\longrightarrow \text{(M-LET, definition of residual)} \\
& \langle r \text{ true} \mid \gamma' \mid \text{residual}(\text{true} :: bs, P) \rangle \\
&\longrightarrow \text{(M-RESUME, } \llbracket r \rrbracket \gamma' = (\sigma, \chi_{\text{count}}(P)) \text{ (log } |\gamma'| = 1 \text{ environment operations))} \\
& \langle \text{return true} \mid \gamma' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(\text{true} :: bs, P) \rangle
\end{aligned}$$

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

1667       **Case**  $1 + |bs| < n$ . We obtain the following configuration.

1668  $\longrightarrow$   $\text{steps}(t)(\text{true}::bs)$  (by Lemma C.4)

1669  $\langle V j \mid \gamma'' \mid (\sigma', \chi_{\text{count}}(P)) :: \text{residual}(\text{true} :: bs, P) \rangle$

1670       where  $?j = \text{labs}(t)(\text{true} :: bs)$ ,  $\gamma'' = \text{env}(t)(\text{true} :: bs)$ ,  $\sigma' = \text{pure}(t)(\text{true} :: bs)$

1671       and  $\llbracket V \rrbracket \gamma'' = (\text{env}^\perp(P), \lambda\_.\mathbf{do}$  Branch  $\langle \rangle)$

1672  $=$  (definition of arrive when  $1 + |bs| < n$ )

1673  $\text{arrive}(\text{true} :: bs, P)$

1674  $\longrightarrow$   $T(\text{true}::bs, n)$  (induction hypothesis)

1675  $\text{depart}(\text{true} :: bs, P)$

1676  $=$  (definition of depart when  $1 + |bs| < n$ )

1677  $\langle \mathbf{return} \ i \mid \gamma \mid \text{residual}(\text{true} :: bs, P) \rangle$

1678       where  $i = c(\text{true} :: \text{true} :: bs) + c(\text{false} :: \text{true} :: bs)$  and  $\gamma = \text{env}_{\text{false}}^\uparrow(\text{true} :: bs, P)$

1679  $=$  (definition of residual and purecont)

1680  $\langle \mathbf{return} \ i \mid \gamma \mid [((\gamma', x_{\text{true}}, \mathbf{let} \ x_{\text{false}} \leftarrow r \ \text{false} \ \mathbf{in} \ x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id})] \rangle$

1681       where  $\gamma' = \text{env}_{\text{true}}^\downarrow(bs, P)$

1682  $\longrightarrow$  (M-RETCNT)

1683  $\langle \mathbf{let} \ x_{\text{false}} \leftarrow r \ \text{false} \ \mathbf{in} \ x_{\text{true}} + x_{\text{false}} \mid \gamma'' \mid [(\text{purecont}(bs, P), \chi_{id})] \rangle$

1684       where  $\gamma'' = \gamma'[x_{\text{true}} \mapsto \llbracket i \rrbracket \gamma']$

1685  $\longrightarrow$  (M-LET)

1686  $\langle r \ \text{false} \mid \gamma'' \mid [((\gamma'', x_{\text{false}}, x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id})] \rangle$

1687  $=$  (definition of purecont and residual)

1688  $\langle r \ \text{false} \mid \gamma'' \mid \text{residual}(\text{false} :: bs, P) \rangle$

1689  $\longrightarrow$  (M-RESUME)

1690  $\langle \mathbf{return} \ \text{false} \mid \gamma'' \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(\text{false} :: bs, P) \rangle$

1691       where  $\sigma = \text{pure}(t)(bs)$

1692  $\longrightarrow$   $\text{steps}(t)(\text{false}::bs)$  (by Lemma C.4 and assumption  $|false :: bs| < n$ )

1693  $\langle V j \mid \gamma \mid (\sigma, \chi_{\text{count}}(P)) :: \text{residual}(\text{false} :: bs, P) \rangle$

1694       where  $?j = \text{labs}(t)(\text{false} :: bs)$ ,  $\sigma = \text{pure}(t)(\text{false} :: bs)$ ,  $\gamma = \text{env}(t)(\text{false} :: bs)$

1695       and  $\llbracket V \rrbracket \gamma = (\text{env}^\perp(P), \lambda\_.\mathbf{do}$  Branch  $\langle \rangle)$

1696  $=$  (definition of arrive when  $1 + |bs| < n$ )

1697  $\text{arrive}(\text{false} :: bs, P)$

1698  $\longrightarrow$   $T(\text{false}::bs, n)$  (induction hypothesis)

1699  $\text{depart}(\text{false} :: bs, P)$

1700  $=$  (definition of depart when  $1 + |bs| < n$ )

1701  $\langle \mathbf{return} \ j \mid \gamma \mid \text{residual}(\text{false} :: bs, P) \rangle$

1702       where  $j = c(\text{true} :: \text{false} :: bs) + c(\text{false} :: \text{false} :: bs)$  and  $\gamma = \text{env}_{\text{false}}^\uparrow(\text{false} :: bs, P)$

1703  $=$  (definition of residual and purecont)

1704  $\langle \mathbf{return} \ j \mid \gamma \mid [((\gamma'', x_{\text{false}}, x_{\text{true}} + x_{\text{false}}) :: \text{purecont}(bs, P), \chi_{id})] \rangle$

1705  $\longrightarrow$  (M-RETCNT)

1706  $\langle x_{\text{true}} + x_{\text{false}} \mid \gamma''[x_{\text{false}} \mapsto \llbracket j \rrbracket \gamma''] \mid \text{residual}(bs, P) \rangle$

1707  $\longrightarrow$  (M-PLUS)

1708  $\langle \mathbf{return} \ m \mid \gamma''[x_{\text{false}} \mapsto \llbracket j \rrbracket \gamma''] \mid \text{residual}(bs, P) \rangle$

1709       where

1710  $m = c(\text{true} :: \text{true} :: bs) + c(\text{false} :: \text{true} :: bs) + c(\text{true} :: \text{false} :: bs) + c(\text{false} :: \text{false} :: bs)$

1711        $= c(\text{true} :: bs) + c(\text{false} :: bs) = c(bs) \leq 2^{n-|bs|}$

1712  $=$  (definition of depart when  $|bs| < n$ )

1713  $\text{depart}(bs, P)$

1714

1715

$$\begin{aligned}
& \text{Step analysis. The total number of machine transitions is given by} \\
& 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_{2 \leq |bs'| \leq n-|bs|} \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}$$



1814 *Step analysis.* The total number of machine transitions is given by

$$\begin{aligned}
1815 & \\
1816 & \\
1817 & \\
1818 & \\
1819 & \quad 9 + \text{steps}(t)(\text{true} :: bs) + T(\text{true} :: bs, n) + \text{steps}(t)(\text{false} :: bs) + T(\text{false} :: bs, n) \\
1820 & = \quad (\text{reorder}) \\
1821 & \quad 9 + T(\text{true} :: bs, n) + T(\text{false} :: bs, n) + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
1822 & = \quad (\text{definition of } T \text{ when } |bs| + 1 = n) \\
1823 & \quad 9 + 2 + 2 + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
1824 & = \quad (\text{simplify}) \\
1825 & \quad 9 + 2^2 + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
1826 & = \quad (\text{rewrite } 2 = n - |bs| + 1) \\
1827 & \quad 9 + 2^{n-|bs|+1} + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
1828 & = \quad (\text{multiply by 1}) \\
1829 & \quad 9 * (2^{n-|bs|} - 1) + 2^{n-|bs|+1} + \text{steps}(t)(\text{true} :: bs) + \text{steps}(t)(\text{false} :: bs) \\
1830 & = \quad (\text{rewrite binary sum}) \\
1831 & \quad 9 * (2^{n-|bs|} - 1) + 2^{n-|bs|} + \sum_{\substack{1 \leq |bs'| \leq n-|bs| \\ bs' \in \mathbb{B}^*}} \text{steps}(t)(bs' ++ bs) \\
1832 & \\
1833 & \\
1834 & = \quad (\text{definition of } T) \\
1835 & \quad T(bs, n) \\
1836 & \\
1837 & \\
1838 & \\
1839 & \\
1840 & \quad \square \\
1841 & \\
1842 & \\
1843 & \\
1844 & \\
1845 &
\end{aligned}$$

1846 The following theorem is a copy of Theorem 5.10.

1847  
1848  
1849 **THEOREM C.6.** *For all  $n > 0$  and any  $n$ -standard predicate  $P$  it holds that*

- 1850  
1851  
1852  
1853 (1) *The program effcount is a generic counting program*  
1854 (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.

1863  
 1864  
 1865  
 1866  
 1867  
 1868  
 1869  
 1870  
 1871  
 1872  
 1873  
 1874  
 1875  
 1876  
 1877  
 1878  
 1879  
 1880  
 1881  
 1882  
 1883      $\langle \text{effcount } P \mid \emptyset \mid [(\square, \chi_{id})] \rangle$   
 1884     = (definition of residual)  
 1885      $\langle \text{effcount } P \mid \emptyset \mid \text{residual}(P, \square, t, c) \rangle$   
 1886      $\longrightarrow$  (M-APP,  $\llbracket \text{effcount} \rrbracket \emptyset = (\emptyset, \lambda \text{pred.} \dots)$ )  
 1887      $\langle \mathbf{handle } \text{pred } (\lambda \_ . \mathbf{do } \text{Branch } \langle \rangle) \mathbf{with } H_{\text{count}} \mid \gamma \mid \text{residual}(P, \square) \rangle$   
 1888     where  $\gamma = \text{env}^\perp(P)$   
 1889      $\longrightarrow$  (M-HANDLE)  
 1890      $\langle \text{pred } (\lambda \_ . \mathbf{do } \text{Branch } \langle \rangle) \mid \gamma \mid (\square, (\gamma, H_{\text{count}})) \text{ :: residual}(P, \square) \rangle$   
 1891     = (definition of  $\chi_{\text{count}}$ )  
 1892      $\langle \text{pred } (\lambda \_ . \mathbf{do } \text{Branch } \langle \rangle) \mid \gamma \mid (\square, \chi_{\text{count}}(P)) \text{ :: residual}(P, \square) \rangle$   
 1893      $\longrightarrow$   $\text{steps}(t)(\square)$  (by Lemma C.4)  
 1894      $\langle (\lambda \_ . \mathbf{do } \text{Branch } \langle \rangle) j \mid \gamma' \mid (\sigma, \chi_{\text{count}}(P)) \text{ :: residual}(P, \square) \rangle$   
 1895     where  $\gamma' = \text{env}(t)(\square)$ ,  $\sigma = \text{pure}(t)(bs)$  and  $?j = \text{labs}(t)(bs)$   
 1896     = (definition of arrive)  
 1897      $\text{arrive}(P, \square)$   
 1898      $\longrightarrow$   $T(\square, n)$  (by Lemma C.5)  
 1899      $\text{depart}(P, \square)$   
 1900     = (definition of depart)  
 1901      $\langle \mathbf{return } m \mid \gamma \mid \text{residual}(P, \square) \rangle$   
 1902     where  $\gamma = \text{env}^\perp(P)$  and  $m = c(\square) \leq 2^{n-|bs|} = 2^n$   
 1903     = (definition of residual)  
 1904      $\langle \mathbf{return } m \mid \gamma \mid [(\square, \chi_{id})] \rangle$   
 1905      $\longrightarrow$  (M-HANDLE-RET,  $H_{id}^{\text{val}} = \{\mathbf{val } x \mapsto \mathbf{return } x\}$ )  
 1906      $\langle \mathbf{return } x \mid \emptyset[x \mapsto m] \mid \square \rangle$   
 1907  
 1908  
 1909  
 1910  
 1911

1912 *Analysis.* The machine yields the value  $m$ . By Lemma C.5 it follows that  $m \leq 2^{n-|bs|} = 2^{n-|\llbracket \rrbracket|} = 2^n$ .  
 1913 Furthermore, the total number of transitions used were

$$\begin{aligned}
 & 5 + \text{steps}(t)(\llbracket \rrbracket) + T(\llbracket \rrbracket, n) \\
 &= \text{(definition of } T) \\
 & 5 + \text{steps}(t)(\llbracket \rrbracket) + 9 * 2^n + 2^{n+1} + \sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') \\
 &= \text{(simplify)} \\
 & 5 + \text{steps}(t)(\llbracket \rrbracket) + 9 * 2^n + 2^{n+1} + \sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') \\
 &= \text{(reorder)} \\
 & 5 + \left( \sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') \right) + \text{steps}(t)(\llbracket \rrbracket) + 9 * 2^n + 2^{n+1} \\
 &= \text{(rewrite as unary sum)} \\
 & 5 + \left( \sum_{bs' \in \mathbb{B}^*}^{1 \leq |bs'| \leq n} \text{steps}(t)(bs') + \sum_{bs' \in \mathbb{B}^*}^{0 \leq |bs'| \leq 0} \text{steps}(t)(bs') \right) + 9 * 2^n + 2^{n+1} \\
 &= \text{(merge sums)} \\
 & 5 + \left( \sum_{bs' \in \mathbb{B}^*}^{0 \leq |bs'| \leq n} \text{steps}(t)(bs') \right) + 9 * 2^n + 2^{n+1} \\
 &= \text{(definition of } \mathcal{O}) \\
 & \left( \sum_{bs' \in \mathbb{B}^*}^{0 \leq |bs'| \leq n} \text{steps}(t)(bs') \right) + \mathcal{O}(2^n)
 \end{aligned}$$

□

## 1940 D PROOF DETAILS FOR THE NO SHORTCUTS LEMMA

1941 The proof of Lemma 5.11 relies on the fact that any  $n$ -standard predicate has a canonical form.  
 1942 Section D.1 disseminates canonical predicates, whilst Section D.2 proves Lemma 5.11.

### 1944 D.1 Canonical Predicates

1945 The decision tree model (Definition 5.2) captures the interaction between a given predicate  $P$  and  
 1946 its point  $p$ . The interior nodes correspond to those places where  $P$  queries  $p$ , whilst the leaves  
 1947 represent answers ultimately conferred from the dialogue between the predicate and its point.

1948 The abstract nature of the decision tree model means that concrete syntactic structure of the  
 1949 predicate is lost. Thus we cannot hope to reconstruct a particular predicate from its model. Indeed  
 1950 many syntactically distinct predicates may share the same model. However, we can construct *some*  
 1951 predicate from a given model, namely, the *canonical predicate*. Intuitively, the canonical predicate  
 1952  $P'$  of  $P$  is a predicate which exhibits the same dialogue as  $P$  for every (valid) point.

1953 Let  $\mathcal{U}(P) := bs \mapsto \mathcal{T}(P)(bs)$ .<sup>1</sup> denote the procedure for constructing an *untimed decision tree* of  
 1954 a given predicate  $P$ .

1956 *Definition D.1 (Canonical predicate).* A canonical predicate  $P'$  of an  $n$ -standard predicate  $P$  is  
 1957 itself an  $n$ -standard predicate whose body (syntactically) consists entirely of **let**-bindings of point  
 1958 applications and whose continuation is either another **let**-expression of the same form or **return**  $b$   
 1959 for some boolean  $b$ . Moreover,  $P'$  exhibits the same dialogue as  $P$ , that is for all  $bs \in \mathbb{B}^*$  such that



1961  $|bs| \leq n$  that

$$1962 \quad \mathcal{U}(P)(bs) = \mathcal{U}(P')(bs)$$

1963 Next we define a procedure for constructing canonical predicate of any given  $n$ -standard predicate.

1964 *Definition D.2 (Normalisation procedure for predicates).* The meta-procedure `norm` takes as input  
 1965 an  $n$ -standard untyped decision tree, and outputs a program whose type is  $\text{Point} \rightarrow \text{Bool}$ , which is  
 1966 exactly the type of predicates. The procedure makes use of an auxiliary procedure `body` to generate  
 1967 the predicate body.

$$1969 \quad \begin{array}{ll} \text{norm} & : (\mathbb{B}^* \rightarrow \text{Lab}) \rightarrow \text{Val} \\ \text{norm}(t) & := \lambda p^{\text{Point}}. \text{body}(t, [], p) \end{array}$$

$$1970 \quad \text{body} : (\mathbb{B}^* \rightarrow \text{Lab}) \times \mathbb{B}^* \times \text{Val} \rightarrow \text{Comp}$$

$$1971 \quad \text{body}(t, bs, p) := \begin{cases} \text{return } b & t(bs) = !b \\ \text{let } b \leftarrow p \text{ in} \\ \text{if } b \text{ then } \text{body}(t, \text{true} :: bs, p) & \text{if } t(bs) = ?i \\ \text{else } \text{body}(t, \text{false} :: bs, p) \end{cases}$$

1972 As convenient notation we write  $\text{norm}(P)$  to mean  $\text{norm}(bs \mapsto \mathcal{U}(P)(bs))$ . Next we show that  
 1973 the meta-procedure `norm` produces canonical predicates.

1974 **LEMMA D.3.** *Suppose  $P$  is an  $n$ -standard predicate then  $P' := \text{norm}(P)$  is an  $n$ -standard predicate  
 1975 such that for all  $bs \in \mathbb{B}^*$ ,  $|bs| \leq n$*

$$1976 \quad \mathcal{U}(P)(bs) = \mathcal{U}(P')(bs')$$

1977 **PROOF.** By induction on  $n$  and `body`.

1978  $\square$

1979 **LEMMA D.4.** *The procedure `norm` generates canonical predicates.*

1980 **PROOF.** First observe that the syntax produced by the `body` procedure of `norm` conforms with the  
 1981 syntactic restrictions of canonical predicates (Definition D.1). The rest follows as by Lemma D.3.  $\square$

## 1982 D.2 No Shortcuts

1983 We now have the necessary machinery to show that every  $n$ -count program in  $\lambda_b$  has at least  
 1984 exponential time complexity. The following lemma is a copy of Lemma 5.11.

1985 **LEMMA D.5.** *If  $C$  is an  $n$ -count program and  $P$  is an  $n$ -standard predicate, then  $C$  applies  $P$  to at  
 1986 least  $2^n$  distinct  $n$ -points. More formally, for any of the  $2^n$  possible semantic  $n$ -points  $\pi : \mathbb{N}_n \rightarrow \mathbb{B}$ ,  
 1987 there is a term  $\mathcal{E}[P p]$  appearing in the small-step reduction of  $C P$  such that  $p$  is a closed value (hence  
 1988 an  $n$ -point) and  $\mathbb{P}[[p]] = \pi$ .*

1989 **PROOF.** Suppose  $C$  and  $P$  are as above, and suppose for contradiction that  $\pi$  is some semantic  
 1990  $n$ -point such that no corresponding application  $P p$  ever arises in the course of computing  $C P$ . Let  
 1991  $t$  be the untyped decision tree for  $P$ . Now consider the leaf node in  $t$  corresponding to the point  $\pi$ ,  
 1992 and let  $t'$  be the tree obtained from  $t$  by simply negating the boolean value at this leaf node, that is

$$1993 \quad t' := bs' \mapsto \begin{cases} \neg b & \text{if } bs = bs' \\ \mathcal{U}(P)(bs') & \text{otherwise} \end{cases}$$

1994 Then  $P' = \text{norm}(t')$  constructs a canonical predicate, and as the numbers of true-leaves in  $t$  and  $t'$   
 1995 differ by 1, it follows that their count at the leaf node in question differ by 1, i.e.

$$1996 \quad |C(P')(bs) - C(P)(bs)| = 1.$$

Taking  $bs = []$ , we get that the values ultimately returned by  $C P$  and  $C P'$  differ by 1, i.e.

$$|C(P')([]) - C(P)([])| = 1.$$

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 have to argue that if the computation of  $C P$  never actually ‘visits’ the leaf node in question, then  $C$  is unable to detect any difference between  $P$  and  $P'$ . To establish our argument we make use of a variation of Milner’s *context lemma* for PCF. Specifically, we have to show the following by induction on length of reduction sequences:

LEMMA D.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  $\text{Nat}$  or  $\text{Bool}$ . If  $\mathcal{F}[P] \rightsquigarrow^m \mathbf{return} V$  then  $\mathcal{F}[P'] \rightsquigarrow^* \mathbf{return} V$  where the type of  $V$  is either  $\text{Nat}$  or  $\text{Bool}$ .*

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

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

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

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

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

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

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  $\text{Point}$ . 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

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

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

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

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

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

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

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

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  $\mathbb{P}[[p]] \neq \pi$ . 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

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

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 D.6 both  $C\ P$  and  $C\ P'$  reduce to the same value which contradicts the initial assumption.

□

2059  
2060  
2061  
2062  
2063  
2064  
2065  
2066  
2067  
2068  
2069  
2070  
2071  
2072  
2073  
2074  
2075  
2076  
2077  
2078  
2079  
2080  
2081  
2082  
2083  
2084  
2085  
2086  
2087  
2088  
2089  
2090  
2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105  
2106  
2107