

Crumbling Abstract Machines

Beniamino Accattoli

LIX

Inria & École Polytechnique

France

beniamino.accattoli@inria.fr

Giulio Guerrieri

Department of Computer Science

University of Bath

United Kingdom

g.guerrieri@bath.ac.uk

Andrea Condoluci

Department of Computer Science and Engineering

University of Bologna

Italy

andrea.condoluci@unibo.it

Claudio Sacerdoti Coen

Department of Computer Science and Engineering

University of Bologna

Italy

claudio.sacerdoticoen@unibo.it

ABSTRACT

Extending the λ -calculus with a construct for sharing, such as let expressions, enables a special representation of terms: iterated applications are decomposed by introducing sharing points in between any two of them, reducing to the case where applications have only values as immediate subterms.

This work studies how such a crumbled representation of terms impacts on the design and the efficiency of abstract machines for call-by-value evaluation. About the design, it removes the need for data structures encoding the evaluation context, such as the applicative stack and the dump, that get encoded in the environment. About efficiency, we show that there is no slowdown, clarifying in particular a point raised by Kennedy, about the potential inefficiency of such a representation.

Moreover, we prove that everything smoothly scales up to the delicate case of open terms, needed to implement proof assistants. Along the way, we also point out that continuation-passing style transformations—that may be alternatives to our representation—do not scale up to the open case.

KEYWORDS

abstract machine, complexity, explicit substitution, lambda-calculus

ACM Reference Format:

Beniamino Accattoli, Andrea Condoluci, Giulio Guerrieri, and Claudio Sacerdoti Coen. 2019. Crumbling Abstract Machines. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 35 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

This paper is about the extension of λ -calculus with explicit constructors for *sharing*. The simplest such construct is a let $x = u$ in t

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Conference'17, July 2017, Washington, DC, USA

© 2019 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

expression, standing for t where x will be substituted by u , that we also write more concisely as $t[x \leftarrow u]$ and call ES (for *explicit sharing*, or *explicit substitution*¹). Thanks to ES, β -reduction can be decomposed into more atomic steps. The simplest decomposition splits β -reduction as $(\lambda x.t)u \rightarrow_{\beta_{ES}} t[x \leftarrow u] \rightarrow_{ES} t\{x \leftarrow u\}$ where $t\{x \leftarrow u\}$ is the meta-level substitution of u for the free occurrences of x in t .

It is well-known that ES are somewhat redundant, as they can always be removed, by simply coding them as β -redexes. They are however more than syntactic sugar, as they provide a simple and yet remarkably effective tool to understand, implement, and program with λ -calculi and functional programming languages.

From a logical point of view, ES are the proof terms corresponding to the extension of natural deduction with a cut rule, and the cut rule is *the* rule representing computation, according to Curry-Howard. From an operational semantics point of view, they allow elegant formulations of subtle strategies such as call-by-need evaluation—various presentations of call-by-need use ES [12, 22, 24, 25, 31, 32] and a particularly simple one is in Accattoli et al. [3]. From a programming point of view, they are part of most functional languages we are aware of. From a rewriting point of view, they enable proof techniques not available within the λ -calculus (e.g. reducing a global rewriting properties such as standardization to a local form, see Accattoli [1]). Finally, sharing is used in all implementations of tools based on the λ -calculus to circumvent *size explosion*, the degenerate behavior for which the size of λ -terms may grow exponentially with the number of β -steps.

Crumbled forms. Once sharing is added to the λ -calculus, it enables a representation of terms where a sharing point is associated with every constructor of the term. Such a special form, roughly, is obtained by (recursively) decomposing iterated applications by introducing an ES in between any two of them. For instance, the representation of the term $((\lambda x.x(xx))y)((\lambda z.z)y)y$ is

$$(w''y)[w'' \leftarrow w'w][w' \leftarrow (\lambda x.(xx'))[x' \leftarrow xx]]y[w \leftarrow (\lambda z.z)y]$$

Note that the transformation involves also function bodies (i.e. $\lambda x.x(xx)$ turns into $\lambda x.(xx')[x' \leftarrow xx]$), that ES are grouped together

¹let expressions and explicit substitutions usually come with different operational semantics: let expressions substitute in just one step, while explicit substitutions substitute in many *micro* steps, percolating through the term structure. They follow however the same typing principles. Moreover, explicit substitutions have many different formulations. In this paper we see let expressions as yet another form of explicit substitutions, and thus conflate the two terminologies.

unless forbidden by abstractions, and that ES are flattened out, i.e. they are not nested unless nesting is forced by abstractions.

This work studies such a representation, called *crumbled* as it crumbles a term by means of ES. Our *crumbling translation* closely resembles—while not being exactly the same—the transformation into *administrative normal form* (shortened to ANF) introduced by Flanagan et al. [17], building on work by Sabry and Felleisen [30], itself a variant of the *continuation-passing style* (CPS) translation.

A delicate point is to preserve crumbled forms during evaluation. ES often come together with commutation rules to move them around the term structure. These rules are often used to unveil redexes during evaluation or to preserve specific syntactic forms. They may introduce significant overhead that, if not handled carefully, can even lead to asymptotic slowdowns as shown by Kennedy [21]. One of the contributions of this work is to show that crumbled forms can be evaluated and preserved with no need of commutation rules, therefore avoiding Kennedy's potential slowdown.

This paper. Our focus is on the impact of crumbled forms on the design and asymptotic overhead of abstract machines with *weak* evaluation (i.e. out of abstractions) on closed terms, and the scalability to (possibly) open terms. Bounding the overhead of abstract machines is a new trend, according to which the machine overhead has to be proved polynomial or even linear in the number of β -steps [2–5, 8, 10]. Open terms—that are not needed to implement functional languages—are used to implement the more general and subtle case of proof assistants. The two topics actually motivate each other: the naive handling of open terms with the techniques for functional languages gives abstract machines with exponential overhead [8, 10], which pushes to develop more efficient machines.

We anticipate here the main results of the paper: crumbled forms induce abstract machines for weak evaluation with less data structures and the transformation does not introduce any asymptotic overhead. Moreover, these facts smoothly scale up to open terms.

Why study crumbled forms. Our interest in studying crumbled forms comes precisely from the fact that they remove some data structures from the design of abstract machines. The relevance of this fact becomes evident when one tries to design abstract machines for strong evaluation (that is, evaluating under abstraction). The study of such machines is extremely technical (see also section Sect. 8) because they have more data structures and more transitions than in the closed and open cases. The many additional transitions are in particular due to the handling of the various data structures. In call-by-name, the situation is still manageable [2, 4, 13, 18], but in call-by-value/need the situation becomes quickly desperate—it is not by chance that there is not a single strong abstract machine for call-by-value/need in the literature.

This work is then preliminary to a detailed study of strong abstract machines for call-by-value and call-by-need. The aim is to explore the subtleties in frameworks that are well understood, such as the closed and open call-by-value cases, and show that there are no slowdowns in turning to a crumbled representation.

The next sub-sections continue the introduction with a lengthy overview of the role of environments, the content of the paper, the relationship with the ANF, the asymptotic study of abstract machines, and related work.

1.1 Environments

ES are often grouped together instead of being scattered all over the term, in finite sequences called *environments*. Abstract machines typically rely on environments. Crumbled forms also rely on packing ES together, as pointed out before, but depart from the ordinary case as environments may appear also under abstractions.

Crumbled Environments. The notion of environment induced by crumbled forms, named here *crumbled environments*, is peculiar. Crumbled environments indeed play a double role: they both store delayed substitutions, as also do ordinary environments, and *encode evaluation contexts*. In ordinary abstract machines, the evaluation context is usually stored in data structures such as the *applicative stack* or the *dump*. Roughly, they implement the search for the redex in the ordinary applicative structure of terms. For crumbled forms, the evaluation context is encoded in the crumbled environment, and so the other structures disappear.

Operations on Crumbled Environments. There are two subtle implementative aspects of crumbled environments, that set them apart from ordinary ones. Ordinary environments are presented with a sequential structure but they are only accessed randomly (that is, not sequentially)—in other words, their sequential structure does not play a role. Crumbled environments, as the ordinary ones, are accessed randomly, to retrieve delayed substitutions, but they are also explored sequentially—since they encode evaluation contexts—in order to search for redexes. Therefore, their implementation has to reflect the sequential structure.

The second subtlety is that crumbled machines also have to concatenate environments, that is an operation never performed by ordinary machines, and that has to be concretely implemented as efficiently as possible, i.e. in constant time. That this point is subtle is proved by the fact that Kennedy's slowdown [21] amounts to a quadratic overhead in evaluating terms in ANF due to the concatenation of environments.

To address these points, we provide a prototype OCaml implementation of crumbled environments in the appendix of [?], to be compared with the one of global environments in Accattoli and Barras [5], that does not concretely implement the sequential structure. In particular, our implementation concatenates environments in constant time and does not suffer from Kennedy's slowdown. Essentially, Kennedy's slowdown amounts to the fact that his implementation concatenates ANF environments in linear rather than constant time (see Section 9).

1.2 Content of the Paper

The Closed Case. First, we define crumbled forms and an abstract machine evaluating them, the Crumble GLAM, and show that it implements Plotkin's *closed* small-step call-by-value (CbV for short) λ -calculus (extended with conditionals, see below). Moreover, we study the overhead of the machine, and show that it is linear in the number of β -steps and in the size of the initial term, exactly as the best machines for CbV executing ordinary terms. Therefore, the crumbling transformation does not introduce any asymptotic overhead. The study is detailed and based on a careful and delicate spelling of the invariants of the machine. In particular, our approach does not suffer from Kennedy's potential slowdown.

Open Terms. The second ingredient of the new trend of abstract machines [2–5, 8, 10]—the first being complexity analyses—is studying evaluation in presence of (possibly) open terms or even strong evaluation (*i.e.* under abstraction), which is required in order to implement proof assistants. Apart from few exceptions—Crégut [13], Grégoire and Leroy [19], and García-Pérez et al. [18]—the literature before the new wave mostly neglected these subtle cases, and none of those three papers addressed complexity.

The open case, in which evaluation is weak but terms are possibly open is strictly harder than the closed one, and close in spirit to the strong case, but easier to study—it is for instance the one studied by Grégoire and Leroy [19] when modeling (an old version of) the abstract machine of the kernel of Coq.

Open Call-by-Value. Open evaluation for CbV—shortened Open CbV—is particularly subtle because, as it is well-known, Plotkin’s operational semantics is not adequate when dealing with open terms—see Accattoli and Guerrieri [7, 9]. Open CbV has been studied deeply by Accattoli and Guerrieri [7, 8, 9], Accattoli and Sacerdoti Coen [10], exploring different presentations, their rewriting, cost models, abstract machines, and denotational semantics. One of the motivations of this work is to add a new piece to the puzzle, by lifting the crumbling technique to the open case.

Our second contribution is to show that the crumbling technique smoothly scales up to Open CbV. We provide an abstract machine, the Open Crumble GLAM, and we show that it implements *the fireball calculus*—the simplest presentation of Open CbV—and that, as in the closed case, it only has a linear overhead. Two aspects of this study are worth pointing out. First, the technical development follows almost identically the one for the closed case, once the subtler invariants of the new machine have been found. Second, *the substitution of abstractions on demand*, a technical optimizations typical of open/strong cases (introduced in Accattoli and Dal Lago [6] and further studied in Accattoli and Guerrieri [8], Accattoli and Sacerdoti Coen [10]), becomes superfluous as it is subsumed by the crumbling transformation.

1.3 The Relationship with ANF

As long as one sticks to the untyped λ -calculus, crumbled forms coincide with ANF. The ANF, we said, is a variant of the CPS transformation. Roughly, the difference is that the ANF does not change the type, when terms are typed (here we work without types).

Kennedy [21] pointed out two problems with the ANF. One is the already discussed quadratic overhead, that does not affect our approach. The second one is the fact that the ANF does not smoothly scale up when the λ -calculus is extended to further constructs such as conditionals or pattern matching. Essentially, the ANF requires conditionals and pattern matching to be out of ES, that is, to never have an expression such as $s[x \leftarrow (\text{if } v \text{ then } t \text{ else } u)]$. Unfortunately, these configurations can be created during evaluation. To preserve the ANF, one is led to add so-called *commuting conversions* such as:

$$s[x \leftarrow (\text{if } v \text{ then } t \text{ else } u)] \rightarrow \text{if } v \text{ then } (s[x \leftarrow t]) \text{ else } (s[x \leftarrow u]) \quad (\text{CC})$$

Clearly, there is an efficiency issue: the commutation causes the duplication of the subterm s . A way out is to use a continuation-like technique, which makes Kennedy conclude that then there is no point in preferring ANF to CPS.

This is where our crumble representation departs from the ANF, as we do not require conditionals and pattern matching to be out of ES. Kennedy only studies the closed case. Our interest in open and strong evaluation is to explore the theory of implementation needed for proof assistants. In these settings, commutations of conditionals and pattern matching such as those hinted at by Kennedy are not valid: they are not validated by dependent type systems like those of Coq or Agda. For example, the CC rule above when the conditional is dependently typed breaks the property of *subject reduction*, as typed terms reduce to ill-typed terms. Consider the term:

$$(x + 1)[x : (\text{if } \text{true} \text{ then } \text{nat} \text{ else } \text{bool}) \leftarrow \text{if } \text{true} \text{ then } 0 \text{ else } \text{false}] : \text{nat}$$

that has type *nat* because the type of x is convertible to *nat*. By applying rule CC, we obtain:

$$\text{if } \text{true} \text{ then } ((x + 1)[x \leftarrow 0]) \text{ else } ((x + 1)[x \leftarrow \text{false}])$$

which is clearly ill-typed (in the underlined part).

The problem in the open case is actually more general, as not even the CPS would work: its properties do not scale up to open terms. In Section 9, indeed, we provide a counter-example to the simulation property in the open case.²

To sum up, neither commuting conversions nor the CPS transformation can be used in our framework. Therefore, we accept that conditionals and pattern matching may appear in ES (in contrast to Kennedy) and so depart from the ANF.

In the paper we treat the cases of the closed and open CbV calculi extended with conditionals. The essence of the study is the crumbling of β -reduction, not the conditionals. Conditionals are included only to stress the difference with respect to the ANF (pattern matching can be handled analogously), but they do not require a special treatment.

1.4 The Complexity of Abstract Machines

Asymptotic Bounds vs Benchmarking. The study of asymptotic bounds for abstract machines is meant to complement the use of benchmarking, by covering *all* possible cases, that certainly cannot be covered via benchmarking.

The relevance of such a study is evident when one considers open terms or strong evaluation. For strong evaluation, for instance, for more than 25 years in the literature there has been only Crégut’s abstract machines [13], which on size exploding families of terms actually has exponential overhead (in the number of β -steps and the size of the initial term). A polynomial machine, developed via a careful asymptotic study, is in Accattoli [2]. Similarly, the abstract machine for open terms described in Grégoire and Leroy [19] suffers of exponential overhead on size exploding families (even if the authors then in practice implement a slightly different machine with polynomial overhead). The asymptotic study of this case is in Accattoli and Guerrieri [8], Accattoli and Sacerdoti Coen [10].

Abstract machines vs compilation. Abstract machines and compilation to machine language are two distinct techniques to execute a program. Compilation is typically more efficient, but it only handles the case where terms are closed and evaluation is weak, that is, the one of functional languages. Strong evaluation is sometimes

²Danvy and Filinski [16] claim that the CPS transformation scales up to open terms (their Theorem 2). However, as we discuss in Section 9, they consider only Plotkin’s operational semantics, which is not adequate for open terms.

employed during compilation to optimize the compiled code, but typically only on linear redexes where size and time explosions are not an issue. Abstract machines are the only execution technique implemented in interactive theorem provers based on dependent types, that need strong evaluation.

Kennedy [21] argues that CPS-based translations are superior to ANF also because the CPS makes join points explicit as continuations, so that invocation of the continuation can be compiled efficiently using jumps. The argument is only valid for compilation and it does not affect abstract machines.

Garbage collection. We study *abstract* machines, which on purpose ignore many details of concrete implementations such as garbage collection, which is an orthogonal topic. In particular, garbage collection is always at most polynomial, if not linear, so its omission does not hide harmful blowups. As far as we know, no abstract machine implemented in interactive theorem provers performs garbage collection.

1.5 Related Work

Environments. In a recent work, Accattoli and Barras [5] compare various kinds of environments, namely, *global*, *local*, and *split*, from implementative and complexity points of view. The crumbling transformation can be studied with respect to every style of environment. Here we focus on crumbled *global* environments because they are simpler and because we also consider the open case, where all kinds of environment induce the same complexity.

Administrative Normal Forms. The literature on ANF is scarce. Beyond the already cited original papers, Danvy has also studied them and their relationship to CPS, but usually calling them *monadic normal forms* [14, 15, 20] because of their relationship with Moggi's monadic λ -calculus [26]. That terminology however sometimes describes a more liberal notion of terms, for instance in Kennedy [21], which is also another relevant piece in the literature on ANF.

All proofs and some supplementary material of our paper are in the Appendix of [?] , the long version of this paper.

2 THE PIF CALCULUS

The grammars and the small-step operational semantics of the *Pif calculus* $\lambda_{\text{Plot}}^{\text{if}}$, that is, Plotkin's calculus λ_{Plot} [28] for Closed CbV evaluation extended with booleans and an if-then-else construct, plus error handling for clashing constructs, are in Fig. 1.

A term is either an application of two terms, an if-then-else, or a *value*, which is in turn either a variable, a (λ) -abstraction, true, false, or an error *err*. We distinguish values that are not variables, noted $v_{\neg x}$ and called *practical values*, following Accattoli and Sacerdoti Coen [11]. The *body* of an abstraction $\lambda x.t$ is t and the *bodies* of a conditional if t then u else s are its two branches u and s . Terms are always identified up to α -equivalence and the set of free variables of a term t is denoted by $\text{fv}(t)$; t is *closed* if $\text{fv}(t) = \emptyset$, *open* otherwise. We use $t\{x \leftarrow u\}$ for the term obtained by the capture-avoiding substitution of u for each free occurrence of x in t .

Contexts. In general, *contexts* are denoted by C and are terms with exactly one occurrence of a special constant $\langle \cdot \rangle$ called the *hole*, that is a placeholder for a removed subterm. In the paper we

TERMS	t, u, s	$::=$	$v \mid tu \mid \text{if } t \text{ then } u \text{ else } s$
VALUES	v	$::=$	$x \mid v_{\neg x}$
PRACTICAL VALUES	$v_{\neg x}$	$::=$	$\lambda x.t \mid \text{true} \mid \text{false} \mid \text{err}$
RIGHT V-CONTEXT	R	$::=$	$\langle \cdot \rangle \mid tR \mid Rv \mid \text{if } R \text{ then } u \text{ else } s$

REDUCTION RULES AT TOP LEVEL			
$(\lambda x.t)v$	\mapsto_{β_v}	$t\{x \leftarrow v\}$	
if true then t else u	\mapsto_{ift}	t	
if false then t else u	\mapsto_{iff}	u	
if t then u else s	\mapsto_{ife}	err	if $t = \lambda x.u$ or $t = \text{err}$
tu	$\mapsto_{@e}$	err	if $t \in \{\text{true}, \text{false}, \text{err}\}$

CONTEXTUAL CLOSURE			
$R\langle t \rangle \rightarrow_a R\langle u \rangle$	if	$t \mapsto_a u$	for $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e\}$
\rightarrow_{pif}	$::=$	$\rightarrow_{\beta_v} \cup \rightarrow_{\text{ift}} \cup \rightarrow_{\text{iff}} \cup \rightarrow_{\text{ife}} \cup \rightarrow_{@e}$	

Figure 1: Pif calculus $\lambda_{\text{Plot}}^{\text{if}}$.

use various notions of contexts in different calculi—for $\lambda_{\text{Plot}}^{\text{if}}$ the relevant notion is *right (evaluation) v-context* R (see Fig. 1). The basic operation on (whatever notion of) contexts is the *plugging* $C\langle t \rangle$ of a term t for the hole $\langle \cdot \rangle$ in C : simply the hole is removed and replaced by t , possibly capturing variables.

Evaluation. According to the definition of right v-context, CbV evaluation \rightarrow_{pif} in $\lambda_{\text{Plot}}^{\text{if}}$ is *weak*, i.e. it does not reduce under λ -abstractions and in the branches of an if-then-else. CbV evaluation is defined for *any* (possibly open) term. But it is well-known that this operational semantics is *adequate* only for *closed* terms, as first noticed by Paolini and Ronchi Della Rocca [29], see also Accattoli and Guerrieri [7, 9, ?] and [?]. When restricted to closed terms, $\lambda_{\text{Plot}}^{\text{if}}$ is called *Closed (Conditional) CbV*: in this setting, evaluation can fire a β -redex $(\lambda x.t)u$ only if the argument u is a closed value, i.e. a closed λ -abstraction, a boolean, or *err*; and in the production Rv for the definition of right v-contexts, v is always a closed value. Note that we work with right-to-left evaluation—this is forced by the production Rv in the definition of right evaluation v-contexts. In the closed case one could as well work with left-to-right evaluation, the choice is inessential.

The error constant *err* is generated during evaluation by the two cases of construct clashes: when the condition for an if-then-else is an abstraction and when a boolean is applied to a term. Both cases would be excluded by typing, but in our untyped setting they are possible, and handled via errors. Similarly, errors are also propagated when they appear as conditions for if-then-else and as left terms of an application. These cases are handled by rules \rightarrow_{ife} and $\rightarrow_{@e}$. Note that errors do not propagate when they occur as arguments of applications: if the left sub-term of the application becomes an abstraction that erases the error then the error is *handled* and it is not observable.

A key property of Plotkin's Closed CbV is *harmony*: a closed term is β_v -normal if and only if it is a (closed) value i.e. a (closed) λ -abstraction. Therefore, every closed term either diverges or it evaluates to a (closed) λ -abstraction. Harmony extends to $\lambda_{\text{Plot}}^{\text{if}}$.

PROPOSITION 2.1 (PIF HARMONY). *Let t be a closed term. t is \rightarrow_{pif} -normal if and only if t is a value.*

3 CRUMBLD EVALUATION, INFORMALLY

Decomposing applications. The idea is to forbid the nesting of non-value constructs such as applications and if-then-else without losing expressive power. To ease the explanation, we focus on nested applications and forget about if-then-else—they do not pose any difficulty. Terms such as $(tu)s$ or $t(us)$ are then represented as $(\lambda x.(xs))(tu)$ and $(\lambda x.(tx))(us)$ where x is a fresh variable. It is usually preferred to use let expressions rather than introducing β -redexes, so that one would rather write $\text{let } x = tu \text{ in } (xs)$ and let $x = us$ in (tx) , or, with ES (aka environment entries),

$$(xs)[x \leftarrow tu] \quad \text{and} \quad (tx)[x \leftarrow us].$$

If the crumbling transformation $\underline{\cdot}$ is applied to the whole term—recursively on t , u and s in our examples—all applications have the form vv' , i.e. they only involve values. If moreover CbV evaluation is adopted, then such a crumbled form is stable by evaluation (reduction steps are naturally defined so that a crumbled form reduces to a crumbled form), as variables can only be replaced by values.

Simulation and no evaluation contexts. Let us now have a look at a slightly bigger example and discuss the recursive part of the crumbling transformation. Let $I = \lambda x.x$ be the identity and consider the term $t := ((\lambda y.yy)I)((II)I)$ whose right-to-left evaluation is

$$\begin{array}{llll} t & \rightarrow_{\beta_v} & ((\lambda y.yy)I)(II) & \rightarrow_{\beta_v} & ((\lambda y.yy)I)I \\ & \rightarrow_{\beta_v} & (II)I & \rightarrow_{\beta_v} & II \quad \rightarrow_{\beta_v} \quad I \end{array}$$

The crumbling transformation decomposes all applications, taking special care of grouping all the environment entries together, flattening them out (that is, avoiding having them nested one into the other), and reflecting the evaluation order in the arrangement of the environment. For instance, the crumbled representation \underline{t} of the term t above is

$$\underline{t} = (wz)[w \leftarrow (\lambda y.yy)I][z \leftarrow xI][x \leftarrow II]$$

and evaluation takes always place at the end of the environment:

$$\begin{array}{llll} \underline{t} & \rightarrow_{\beta_v} & (wz)[w \leftarrow (\lambda y.yy)I][z \leftarrow xI][x \leftarrow I] \\ & \rightarrow_{[]} & (wz)[w \leftarrow (\lambda y.yy)I][z \leftarrow II] \\ & \rightarrow_{\beta_v} & (wz)[w \leftarrow (\lambda y.yy)I][z \leftarrow I] \\ & \rightarrow_{[]} & (wI)[w \leftarrow (\lambda y.yy)I] \rightarrow_{\beta_v} (wI)[w \leftarrow II] \\ & \rightarrow_{\beta_v} & (wI)[w \leftarrow I] \rightarrow_{[]} II \rightarrow_{\beta_v} I \end{array}$$

where the \rightarrow_{β_v} steps correspond exactly to steps in the ordinary evaluation of t and $\rightarrow_{[]}$ steps simply eliminate the explicit substitution when its content is a value. Note how the transformation makes the redex always appear at the end of the environment, so that the need for searching for it—together with the notion of evaluation context—disappears.

Let us also introduce some terminology. Values and applications of values are *bites*. The transformation, called *crumbling translation*, turns a term into a pair, called *crumble*, of a bite and an environment.

Turning to micro-step evaluation. The previous example covers what happens when the crumbling transformation is paired with small-step evaluation. Abstract machines, however, employ a finer mechanism that we like to call *micro-step* evaluation, where the substitutions due to β -redexes are delayed and represented as new environment entries, and moreover substitution is decomposed as to act on one variable occurrence at a time. In particular, such

a more parsimonious evaluation never removes environment entries because they might be useful later on—garbage collection is assumed to be an orthogonal and independent process. To give an idea of how micro steps work, let's focus on the evaluation of the subterm $(wz)[w \leftarrow (\lambda y.yy)I]$ of our example (because micro-step evaluations are long and tedious), that proceeds as follows:

$$\begin{array}{llll} (wz)[w \leftarrow (\lambda y.yy)I] & \rightarrow_{\beta_v} & (wz)[w \leftarrow yy][y \leftarrow I] & \rightarrow_{[]} \\ (wz)[w \leftarrow yy][y \leftarrow I] & \rightarrow_{[]} & (wz)[w \leftarrow II][y \leftarrow I] & \rightarrow_{\beta_v} \\ (wz)[w \leftarrow x][x \leftarrow I][y \leftarrow I] & \rightarrow_{[]} & (wz)[w \leftarrow I][x \leftarrow I][y \leftarrow I] & \rightarrow_{[]} \\ (Iz)[w \leftarrow I][x \leftarrow I][y \leftarrow I] & & & \end{array}$$

where \rightarrow_{β_v} steps now introduce new environment entries. Now the redex is not always at the end of the environment, but it is always followed on the right by an environment whose entries are all abstractions, so that the search for the next redex becomes a straightforward visit from right to left of the environment—the evaluation context has been coded inside the sequential structure of the environment.

Abstraction bodies and the concatenation of environments. There is a last point to explain. We adopt weak evaluation—that only evaluates out of abstractions—but the crumbling transformation also transforms the bodies of abstractions and the branches of if-then-else into crumbles. Let us see another example. The crumbled representation of $u := (\lambda x.((xx)(xx)))(II)$ then is

$$\underline{u} = ((\lambda x.((yz)[y \leftarrow xx][z \leftarrow xx]))w)[w \leftarrow II]$$

Micro-step evaluation goes as follows:

$$\begin{array}{ll} \underline{u} & \rightarrow_{\beta_v} ((\lambda x.((yz)[y \leftarrow xx][z \leftarrow xx]))w)[w \leftarrow w'] [w' \leftarrow I] \\ & \rightarrow_{[]} ((\lambda x.((yz)[y \leftarrow xx][z \leftarrow xx]))w)[w \leftarrow I] [w' \leftarrow I] \\ & \rightarrow_{[]} ((\lambda x.((yz)[y \leftarrow xx][z \leftarrow xx]))I)[w \leftarrow I] [w' \leftarrow I]. \end{array}$$

At this point, the reduction of the β -redex (involving λx) has to combine the crumble of the redex itself with the one of the body of the abstraction, by concatenating the environment of the former (here $[w \leftarrow I][w' \leftarrow I]$) at the end of the environment of the latter $([y \leftarrow xx][z \leftarrow xx])$, interposing the entry created by the redex itself $([x \leftarrow I])$, thus producing the new crumble:

$$(yz)[y \leftarrow xx][z \leftarrow xx][x \leftarrow I][w \leftarrow I][w' \leftarrow I].$$

The key conclusion is that evaluation needs to *concatenate* crumbled environments, which is an operation that ordinary abstract machines instead never perform.

Note that transforming abstraction bodies may produce nested ES, if the abstraction occurs in an ES. This is the only kind of nesting of ES that is allowed.

4 THE CRUMBLING TRANSFORMATION

In this section we formally define the language of crumbled forms and the crumbling transformation.

Crumbled forms. Terms are replaced by *crumbles*, which are formed by a bite and an environment, where in turn

- a *bite* is either a *crumbled value* (i.e. a variable, a boolean, an error, or an abstraction over a crumble), an application of crumbled values, or a if-then-else on a crumbled value whose alternatives are crumbles, and
- an *environment* is a finite sequence of explicit substitutions of bites for variables.

Formally, the definition is by mutual induction:

BITES	$b, b' ::= v \mid vw \mid \text{if } v \text{ then } c \text{ else } d$
CRUMBLD VALUES	$v, w ::= x \mid \lambda x.c \mid \text{true} \mid \text{false} \mid \text{err}$
ENVIRONMENTS	$e, e' ::= \epsilon \mid e[x \leftarrow b]$
CRUMBLES	$c, d ::= (b, e)$

Bodies: the bodies of abstractions and if-then-else are themselves crumbles—the forthcoming crumbling transformation is indeed *strong*, as it also transforms bodies.

Crumbles are not closures: the definition of crumbles may remind one of *closures* in abstract machines with local environments, but the two concepts are different. The environment e of a crumble (b, e) , indeed, does not in general bind all the free variables of the bite b .

We freely consider environments as lists extendable on both ends, and whose concatenation is obtained by simple juxtaposition. Given a crumble (b, e) and an environment e' the *appending* of e' to (b, e) is $(b, e) @ e' := (b, ee')$.

Free variables, α -renaming, and all that. All syntactic expressions are *not* considered up to α -equivalence. Free variables are defined as expected for bites. For environments and crumbles they are defined as follows (via the auxiliary notion of *domain* of environments; this is because global environments are used here):

$$\begin{aligned} \text{dom}(e[x \leftarrow b]) &:= \text{dom}(e) \cup \{x\} & \text{dom}(\epsilon) &:= \emptyset \\ \text{dom}((b, e)) &:= \text{dom}(e) & \text{fv}(\epsilon) &:= \emptyset \\ \text{fv}(e[x \leftarrow b]) &:= (\text{fv}(e) \setminus \{x\}) \cup \text{fv}(b) \\ \text{fv}((b, e)) &:= (\text{fv}(b) \setminus \text{dom}(e)) \cup \text{fv}(e). \end{aligned}$$

Let $e = [x_1 \leftarrow b_1] \dots [x_k \leftarrow b_k]$ be an environment: we denote the lookup of x_i in e by $e(x_i) := b_i$. We say that a crumble c or an environment e are *well-named* if all the variables occurring on the lhs of ES outside abstractions in c or e are pairwise distinct.

The crumbling translation. A term is turned into a crumble via the following *crumbling translation* $\underline{\cdot}$, which uses an auxiliary translation $\bar{\cdot}$ from values into crumbled values.

$$\bar{x} := x \quad \overline{\lambda x.t} := \lambda x.\bar{t} \quad \overline{\text{true}} := \text{true} \quad \overline{\text{false}} := \text{false} \quad \overline{\text{err}} := \text{err}$$

$$\begin{aligned} \underline{v} &:= (\bar{v}, \epsilon) & \underline{vv'} &:= (\bar{v}\bar{v'}, \epsilon) \\ \underline{tv} &:= (x\bar{v}, [x \leftarrow b]e) & (*) \\ \underline{ut} &:= \underline{ux} @ ([x \leftarrow b]e) & (*) \end{aligned}$$

$$\underline{\text{if } v \text{ then } u \text{ else } s} := (\text{if } \bar{v} \text{ then } \underline{u} \text{ else } \underline{s}, \epsilon)$$

$$\underline{\text{if } t \text{ then } u \text{ else } s} := (\text{if } x \text{ then } \underline{u} \text{ else } \underline{s}, [x \leftarrow b]e) \quad (*)$$

(*) if t is not a value and $\bar{t} = (b, e)$, and x is fresh.

According to the definition, if u and t are not values, $\underline{ut} = (yx, [y \leftarrow b']e'[x \leftarrow e])$ with $\bar{t} = (b, e)$, $\bar{u} = (b', e')$ and x, y fresh.

Example 4.1. Let $\delta := \lambda x.xx$ and $I := \lambda x.x$: thus, $\bar{I} = \lambda x.\bar{x} = \lambda x.(x, \epsilon)$ and $\bar{\delta} = \lambda x.\bar{xx} = \lambda x.(xx, \epsilon)$ (since $\underline{xx} = (xx, \epsilon)$) and $\underline{\delta\delta} = (\bar{\delta}\bar{\delta}, \epsilon)$. Therefore,

$$\begin{aligned} \underline{\delta\delta I} &= (z\bar{I}, [z \leftarrow \bar{\delta}\bar{\delta}]) \\ &= (z\lambda x.(x, \epsilon), [z \leftarrow (\lambda x.(xx, \epsilon))\lambda x.(xx, \epsilon)]) \end{aligned}$$

$$\begin{aligned} \underline{\delta\delta(xx)} &= (zw, [z \leftarrow \bar{\delta}\bar{\delta}][w \leftarrow \underline{xx}]) \\ &= (zw, [z \leftarrow (\lambda x.(xx, \epsilon))\lambda x.(xx, \epsilon)][w \leftarrow \underline{xx}]). \end{aligned}$$

The crumbling translation $\underline{\cdot}$ is not surjective: the crumble $c := (xx, [x \leftarrow y])$ is such that $\bar{t} \neq c$ for any term t .

Read back. There is a left inverse for the crumbling translation, called *read-back* and defined by:

$$\begin{aligned} x_{\downarrow} &:= x & (\lambda x.c)_{\downarrow} &:= \lambda x.c_{\downarrow} \\ \text{true}_{\downarrow} &:= \text{true} & \text{false}_{\downarrow} &:= \text{false} \\ \text{err}_{\downarrow} &:= \text{err} & (vw)_{\downarrow} &:= v_{\downarrow}w_{\downarrow} \\ (\text{if } v \text{ then } c \text{ else } d)_{\downarrow} &:= \text{if } v_{\downarrow} \text{ then } c_{\downarrow} \text{ else } d_{\downarrow} \\ (b, e[x \leftarrow b'])_{\downarrow} &:= (b, e)_{\downarrow} \{x \leftarrow b'_{\downarrow}\} & (b, \epsilon)_{\downarrow} &:= b_{\downarrow} \end{aligned}$$

PROPOSITION 4.2 (READ-BACK AND THE CRUMBLING TRANSLATION). *For every term t and every value v , one has $\bar{t}_{\downarrow} = t$ and $\bar{v}_{\downarrow} = v$.*

Remark 4.1 (Crumbling translation, free variables).

- (1) For any term t and any value v , one has $\text{fv}(t) = \text{fv}(\bar{t})$ and $\text{fv}(v) = \text{fv}(\bar{v})$; in particular, t is closed if and only if \bar{t} is so.
- (2) For any bite b and crumble c , $\text{fv}(b_{\downarrow}) = \text{fv}(b)$ and $\text{fv}(c_{\downarrow}) = \text{fv}(c)$.
- (3) The crumbling translation commutes with the renaming of free variables.
- (4) The crumbling translation and the read-back map values to values.

Crumbled contexts. For crumbled forms, we need *contexts* both for environments and crumbles:

$$\begin{aligned} \text{ENVIRONMENT CONTEXTS} & \quad E := e[x \leftarrow \langle \cdot \rangle] \\ \text{CRUMBLE CONTEXTS} & \quad C := \langle \cdot \rangle \mid (b, E). \end{aligned}$$

Crumbles can be plugged into both notions of contexts. Let us point out that the following definition of plugging is slightly unusual as it does a little bit more than just replacing the hole, because simply replacing would not provide a well-formed syntactic object: plugging indeed extracts the environment from the plugged crumble and concatenates it with the environment of the context. Such an unusual operation—that may seem *ad-hoc*—is actually one of the key technical points in order to obtain a clean proof of the implementation theorem (see Section 5.2).

Definition 4.3 (Plugging in crumbled contexts). Let $E = e[x \leftarrow \langle \cdot \rangle]$ be an environment context, C be a crumble context, and $c = (b', e')$ be a crumble. The *plugging* $E\langle c \rangle$ of c in E and the *plugging* $C\langle c \rangle$ of c in C are defined by

$$(e[x \leftarrow \langle \cdot \rangle])\langle (b', e') \rangle := e[x \leftarrow b']e' \quad \langle \cdot \rangle\langle c \rangle := c \quad (b, E)\langle c \rangle := (b, E\langle c \rangle)$$

Example 4.4. In Example 4.1 we have seen that $\underline{\delta\delta I} = (z\lambda x.x_{\epsilon}, [z \leftarrow (\lambda x.(xx)_{\epsilon})\lambda x.(xx)_{\epsilon}])$, where we set $b_{\epsilon} := (b, \epsilon)$ for any bite b . We have that $\underline{\delta\delta I} = C\langle c \rangle$ with $C := (z\lambda x.x_{\epsilon}, [z \leftarrow \langle \cdot \rangle])$ and $c := ((\lambda x.(xx)_{\epsilon})\lambda x.(xx)_{\epsilon}, \epsilon)$.

The notions of well-named, $\text{fv}(\cdot)$, and $\text{dom}(\cdot)$ can be naturally extended to crumble contexts. The definition of read back is extended to crumble contexts by setting $\langle \cdot \rangle_{\downarrow} := \langle \cdot \rangle$ and $(b, e[x \leftarrow \langle \cdot \rangle])_{\downarrow} := (b, e)_{\downarrow} \{x \leftarrow \langle \cdot \rangle\}$. Note however that the unfolding of a crumble context is not necessarily a context, because the hole can be duplicated or erased by the unfolding. For instance, let $C := (x\ x, [x \leftarrow \langle \cdot \rangle])$. Then $C_{\downarrow} = \langle \cdot \rangle \langle \cdot \rangle$ is not a context.

Lemma 4.5 provides the properties of the translation needed to prove the invariants of machines in the next sections.

LEMMA 4.5 (PROPERTIES OF CRUMBLING). *For every term t :*

- (1) Freshness: t is well-named.
- (2) Closure: if t is closed, then $\text{fv}(t) = \emptyset$.
- (3) Disjointedness: $\text{dom}(C) \cap \text{fv}(b) = \emptyset$ if $t = C\langle(b, e)\rangle$.
- (4) Bodies: every body in t is the translation of a term.
- (5) Contextual decoding: if $t = C\langle c \rangle$, then C_{\downarrow} is a right v -context.

5 THE CLOSED CASE

Here we show how to evaluate crumbled forms with a micro-step operational semantics. We build over the work of Accattoli and co-authors, who employ the following terminology:

- *Calculus*: for a small-step semantics where both substitution and search for the redex are meta-level operations;
- *Linear calculus*: for a micro-step semantics where substitution is decomposed—the calculus has ES and possibly a notion of environment if the ES are grouped together—but the search for the redex is still meta-level and expressed via evaluation contexts;
- *Abstract machine*: for a micro-step semantics where both substitution and search for the redex are decomposed. The search for redexes is handled via one or more stacks called *applicative stack*, *dump*, *frame*, and so on; the management of names is also explicit, i.e. not up-to α -equivalence.

The crumbling transformation blurs the distinction between a linear calculus and an abstract machine because it allows using the sequential structure of the environment as the only stack needed to search for redexes.

The operational semantics for crumbled forms we present next is in the style of a linear calculus, because spelling out the straightforward search for redexes is not really informative. Nonetheless, we do call it an *abstract machine*, because of the blurred distinction in the crumble case and because we manage names explicitly. In Section 7 we sketch the actual abstract machine (details are in [?]).

5.1 The Crumble GLAM

Transitions. To introduce the Crumble GLAM (GLAM stands for Global Leroy Abstract Machine) we need some definitions. First, environments and crumbles made out of practical values only are defined and noted as follows:

$$\begin{array}{ll} v\text{-ENVIRONMENTS} & e_v ::= e \mid e_v[x \leftarrow v_{\neg x}] \\ v\text{-CRUMBLES} & c_v ::= (v_{\neg x}, e_v) \end{array}$$

Essentially, a v -environment stands for the already evaluated coda of the environment described in the paragraph about micro-steps in Sect. 3, while v -crumbles are fully evaluated crumbles (i.e. final states of the machine), as we show below.

Second, given a crumble c we use c^α for a crumble obtained by α -renaming the names in the domain of c with fresh ones so that c^α is well-named.

The transitions act on crumbles whose environments are v -environments. The top level transitions are:

$$\begin{aligned} ((\lambda x.c) v, e_v) &\mapsto_{\beta_v} (c @ [x \leftarrow v])^\alpha @ e_v \\ (\text{if true then } c \text{ else } d, e_v) &\mapsto_{\text{ift}} c @ e_v \\ (\text{if false then } c \text{ else } d, e_v) &\mapsto_{\text{iff}} d @ e_v \\ (\text{if } v \text{ then } c \text{ else } d, e_v) &\mapsto_{\text{ife}} (\text{err}, e_v) \\ (vw, e_v) &\mapsto_{@e} (\text{err}, e_v) \\ (x, e_v) &\mapsto_{\text{sub}_{var}} (e_v(x), e_v) \\ (xv, e_v) &\mapsto_{\text{sub}_l} (e_v(x) v, e_v) \\ (\text{if } x \text{ then } c \text{ else } d, e_v) &\mapsto_{\text{sub}_{if}} (\text{if } e_v(x) \text{ then } c \text{ else } d, e_v) \end{aligned} \tag{1} \tag{2} \tag{3} \tag{3} \tag{3}$$

- (1) if $v = \lambda x.e$ or $v = \text{err}$
- (2) if $v \in \{\text{true}, \text{false}, \text{err}\}$
- (3) if $x \in \text{dom}(e_v)$

Transitions are then closed by crumble contexts: for every $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e, \text{sub}_{var}, \text{sub}_l, \text{sub}_{if}\}$ define $C\langle c \rangle \rightarrow_a C\langle d \rangle$ if $c \mapsto_a d$. The transition relation \rightarrow_{Cr} of the Crumble GLAM is defined as the union of all these rules. Let us explain each transition:

- \rightarrow_{β_v} : (forget about the α -renaming for the moment—see the next paragraph) the rule removes a β -redex and introduces an ES $[x \leftarrow v]$ instead of performing the meta-level substitution. Moreover, the environment of the body c of the abstraction and the external environment e_v are concatenated (via the appending operation $@$) interposing $[x \leftarrow v]$.
- *Conditional and error transitions* $\rightarrow_{\text{ift}}, \rightarrow_{\text{iff}}, \rightarrow_{\text{ife}}, \rightarrow_{@e}$: these transitions simply mimics the analogous rules on the Pif calculus, with no surprises.
- *Substitution transitions* $\rightarrow_{\text{sub}_l}, \rightarrow_{\text{sub}_r}, \rightarrow_{\text{sub}_{if}}$: the variable x is substituted by the corresponding crumbled value in the environment e_v , if any. In the closed case, a forthcoming invariant guarantees that $e_v(x)$ is always defined so that side-condition (3) is actually always satisfied. There are no rules to substitute on the right of an application (see below).

According to the definitions of plugging and top level transitions, the transition relation follows *right-to-left* evaluation, as the environment on the right of a redex is a v -environment (made of practical values only), which means that it has already been evaluated (see the harmony property for Crumble GLAM in Prop. 5.3 below). Adopting right-to-left evaluation implies that the Crumble GLAM does not need a rule $\rightarrow_{\text{sub}_r}$ symmetrical to $\rightarrow_{\text{sub}_l}$, whose top level shape would be $(vx, e_v) \mapsto_{e_r} (v e_v(x), e_v)$ with $x \in \text{dom}(e_v)$: indeed, if v is a variable then $\rightarrow_{\text{sub}_l}$ applies to the same redex (vx, e_v) , otherwise v is an abstraction and \rightarrow_{β_v} applies to (vx, e_v) .

The cost and the place of α -renaming. Abstract machines with global environments have to α -rename at some point, this is standard³. In our implementation, renaming is implemented as a copy function. And the cost of renaming is under control because of forthcoming invariants of the machine. This is all standard [5].

³Local environments do allow to avoid renamings, but the simplification is an illusion, as the price is paid elsewhere—see Accattoli and Barras [5]—there is no real way out.

Often the burden of renaming/copying is put on the substitution rules. It is less standard to put it on the β_v -transition, as we do here, but nothing changes. Last, a technical remark: in rule \rightarrow_{β_v} the α -renaming at top level has to pick names that are fresh also with respect to the crumble context enclosing it. This point may seem odd but it is necessary to avoid name clashes, and it is trivially obtained in our concrete implementation, where variable names are memory locations and picking a fresh name amounts to allocating a new location, that is of course new globally.

Definition 5.1 (Reachable crumble). A crumble is *reachable* (by the Crumble GLAM) if it is obtained by a sequence of transitions starting from the translation \underline{t} of a closed term t .

Unchaining abstractions. The substitution performed by the rule $\rightarrow_{\text{sub}_{\text{var}}}$ may seem an unneeded optimization; quite the opposite, it fixes an issue causing quadratic overhead in the machine. The culprits are malicious *chains of renamings*, i.e. environments of the form $[x_1 \leftarrow x_2][x_2 \leftarrow x_3] \cdots [x_n \leftarrow \lambda y.c]$ substituting variables for variables and finally leading to an abstraction. Accattoli and Sacerdoti Coen [10] showed that the key to linear overhead is to perform substitution steps while going through the chain from right to left.

Example 5.2. Consider the crumble $\underline{\delta\delta} = (\bar{\delta}\bar{\delta}, \epsilon)$, where $\bar{\delta} = \lambda x.(xx, \epsilon)$; then:

$$\begin{aligned} \underline{\delta\delta} &\rightarrow_{\beta_v} (xx, [x \leftarrow \bar{\delta}]) \rightarrow_{\text{sub}_I} (\bar{\delta}x, [x \leftarrow \bar{\delta}]) \\ &\rightarrow_{\beta_v} (yy, [y \leftarrow x][x \leftarrow \bar{\delta}]) \rightarrow_{\text{sub}_{\text{var}}} (yy, [y \leftarrow \bar{\delta}][x \leftarrow \bar{\delta}]) \rightarrow_{\text{sub}_I} \dots \end{aligned}$$

In Ex. 4.1 we introduced the crumble $\underline{\delta\delta I} = (z\bar{I}, [z \leftarrow \bar{\delta}\bar{\delta}])$ where $\bar{I} = (\lambda x.(x, \epsilon))$; in accordance with the crumble decomposition shown in Ex. 4.4, we have:

$$\begin{aligned} \underline{\delta\delta I} &\rightarrow_{\beta_v} (z\bar{I}, [z \leftarrow xx][x \leftarrow \bar{\delta}]) \rightarrow_{\text{sub}_I} (z\bar{I}, [z \leftarrow \bar{\delta}x][x \leftarrow \bar{\delta}]) \\ &\rightarrow_{\beta_v} (z\bar{I}, [z \leftarrow yy][y \leftarrow x][x \leftarrow \bar{\delta}]) \\ &\rightarrow_{\text{sub}_I} (z\bar{I}, [z \leftarrow yy][y \leftarrow \bar{\delta}][x \leftarrow \bar{\delta}]) \rightarrow_{\text{sub}_I} \dots \end{aligned}$$

Consider now the *open* crumble

$$c := \underline{\delta\delta(xx)} = (zw, [z \leftarrow \bar{\delta}\bar{\delta}][w \leftarrow xx]).$$

The crumble c is normal because its only possible decomposition of the form $C\langle(b, e_v)\rangle$ is for $e_v = \epsilon$ (as xx is not a practical value), and no transitions apply to the rightmost entry $[w \leftarrow xx]$ since x is free.

The Crumble GLAM satisfies a harmony property.

PROPOSITION 5.3 (HARMONY FOR THE CRUMBLE GLAM). A closed crumble c is normal if and only if it is a v -crumble.

5.2 The Implementation Theorem

To show that the Crumble GLAM correctly implements the Pif calculus, we apply an abstract approach introduced by Accattoli and Guerrieri [8], which we reuse as well in the following sections for other crumble abstract machines and other evaluation strategies of the λ -calculus.

The implementation theorem, abstractly. In Accattoli and Guerrieri [8] it is proven that, given

- a generic abstract machine M , which is a *transitions relation* \rightsquigarrow_M over a set of *states* that splits into

- *principal transitions* \rightsquigarrow_p , that corresponds to the evaluation steps on the calculus, and
- *overhead transitions* \rightsquigarrow_o , that are specific of the machine,
- an evaluation strategy \rightarrow in the λ -calculus, and
- a decoding $(\cdot)_\downarrow$ of states of M into terms,

M correctly implements \rightarrow via $(\cdot)_\downarrow$ whenever $(M, \rightarrow, (\cdot)_\downarrow)$ forms an *implementation system*, i.e. whenever the following conditions are fulfilled (where s and s' stand for generic states of M):

- (1) *Initialization*: there is an encoding $_$ of terms such that $\underline{t}_\downarrow = t$;
- (2) *Principal projection*: $s \rightsquigarrow_p s'$ implies $s_\downarrow \rightarrow s'_\downarrow$;
- (3) *Overhead transparency*: $s \rightsquigarrow_o s'$ implies $s_\downarrow = s'_\downarrow$;
- (4) *Determinism*: \rightsquigarrow_M is deterministic;
- (5) *Halt*: M final states (to which no transition applies) decode to \rightarrow -normal terms;
- (6) *Overhead Termination*: \rightsquigarrow_o terminates.

Our notion of implementation, tuned towards complexity analyses, requires a perfect match between the number of steps of the strategy and the number of principal transitions of the execution.

THEOREM 5.4 (MACHINE IMPLEMENTATION, [8]). If a machine M , a strategy \rightarrow on λ -terms and a decoding \cdot_\downarrow form an implementation system then:

- (1) Executions to derivations: for any M -execution $\rho: \underline{t} \rightsquigarrow_M^* s$ there is a \rightarrow -derivation $d: t \rightarrow^* s_\downarrow$.
- (2) Derivations to executions: for every \rightarrow -derivation $d: t \rightarrow^* u$ there is an M -execution $\rho: \underline{t} \rightsquigarrow_M^* s$ such that $s_\downarrow = u$.
- (3) Principal matching: in both previous points the number $|\rho|_p$ of principal transitions in ρ is exactly the length $|d|$ of the derivation d , i.e. $|d| = |\rho|_p$.

The crumbling implementation system. The states of the Crumble GLAM are crumbles. Its principal transitions are those labeled with $\{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e\}$, while the overhead transitions are those labeled with $\{\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}\}$. We can now show that the Crumble GLAM, Pif evaluation \rightarrow_{pif} and the read-back $(\cdot)_\downarrow$ form an implementation system, that is, that the Crumble GLAM implements the Pif calculus.

We are going to provide five of the six sufficient conditions required by the implementation theorem (Thm. 5.4); the sixth one, the termination of overhead transitions, is subsumed by the finer complexity analysis in Subsect. 5.3.

The sufficient conditions, as usual, are proved by means of a few invariants of the machine, given by Lemma 5.5 below. These invariants are essentially the properties of the translation in Lemma 4.5 extended to all reachable crumbles. One of them—namely *contextual decoding*—however, is weaker because reachable crumbles do not necessarily have the same nice structure as the initial crumbles obtained by translation of a term, as the next remark explains.

Remark 5.1. Even though not all crumble contexts unfold to contexts, crumble contexts obtained by decomposing crumbles given by the translation of terms do (Lemma 4.5.5)—this is the contextual decoding property. Unfortunately, it is not preserved by evaluation. Consider the crumble $c := (\lambda x.x(xx))I = ((\lambda x.(xy, [y \leftarrow xx]))\bar{I}, \epsilon)$ with $\bar{I} = \lambda z.(z, \epsilon)$. Clearly, $c = \langle(\lambda x.x(xx))I\rangle$ where $\langle\cdot\rangle_\downarrow = \langle\cdot\rangle$ is a context. After one β_v step, the crumble c reaches

$(xy, [y \leftarrow xx][x \leftarrow \bar{I}]) = C(\bar{I}, \epsilon)$ for $C := (xy, [y \leftarrow xx][x \leftarrow \langle \cdot \rangle])$. But C unfolds to $C_{\downarrow} = \langle \cdot \rangle \langle \langle \cdot \rangle \langle \cdot \rangle \rangle$, which is not a λ -context.

Proof
p. 19

LEMMA 5.5 (INVARIANTS FOR THE CRUMBLE GLAM). *For every reachable crumble c in the Crumble GLAM:*

- (1) Freshness: c is well-named.
- (2) Closure: $\text{fv}(c) = \emptyset$.
- (3) Bodies: every body occurring in c is a subterm (up to renaming) of the initial crumble.
- (4) Weak contextual decoding: for every decomposition $C(\langle b, e_v \rangle)$ where b is not a crumbled value, if C'' is a prefix of C then C''_{\downarrow} is a right v -context.

Freshness and closure are invariants needed to ensure the basic functioning of the machine. The bodies invariant corresponds to what is often called *subterm invariant*: it is the key invariant for complexity analyses, as it allows to bound the size of duplicated subterms (that are always abstractions) using the size of the initial term. Usually, it is only needed for complexity analyses, while here it is needed for the implementation theorem too (namely, only for the proof of the weak contextual decoding invariant). The weak contextual decoding invariant is crucial to show that principal transitions of the Crumble GLAM project on evaluation steps in $\lambda_{\text{Plot}}^{\text{if}}$.

Proof
p. 21

THEOREM 5.6 (IMPLEMENTATION). *Let c be a crumble reachable by the Crumble GLAM.*

- (1) Initialization: $t_{\downarrow} = t$ for every term t .
- (2) Principal projection: if $c \rightarrow_a d$ then $c_{\downarrow} \rightarrow_a d_{\downarrow}$, for any rule $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e\}$.
- (3) Overhead transparency: if $c \rightarrow_a d$ then $c_{\downarrow} = d_{\downarrow}$ for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}\}$.
- (4) Determinism: the transition \rightarrow_{C_r} is deterministic.
- (5) Halt: if c is \rightarrow_{C_r} -normal then c_{\downarrow} is \rightarrow_{pif} -normal.
- (6) Overhead termination: \rightarrow_a terminates, for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}\}$.

Therefore, the Crumble GLAM, Pif evaluation \rightarrow_{pif} , and the read-back $(\cdot)_{\downarrow}$ form an implementation system.

5.3 Complexity for the closed case

To estimate the cost of the Crumble GLAM, we provide first an upper bound on the number of overhead transitions—namely the substitution ones sub_{var} , sub_I , and sub_{if} —in an execution ρ as a function of the number $|\rho|_p$ of principal transitions. Then we discuss the cost of implementing single transitions. Last, by composing the two analyses we obtain the total cost, that is linear in the number of principal transitions and in the size of the initial term/crumble, that is, the machine is bilinear.

Number of transitions: non-renaming substitutions. Let $\rho : c_0 \rightarrow_{C_r}^* c$ be an execution (i.e. a sequence of transitions) in the Crumble GLAM and let $|\rho|_p, |\rho|_{\text{sub}_{\text{var}}}, |\rho|_{\text{sub}_I}, |\rho|_{\text{sub}_{\text{if}}}$ be the number of principal, sub_{var} , sub_I , and sub_{if} transitions in ρ , respectively. Clearly, a sub_I transition can only be immediately followed by a β_v or a $@e$ transition (since \rightarrow_{C_r} is deterministic), and so $|\rho|_{\text{sub}_I} \leq |\rho|_{\beta_v} + |\rho|_{@e} + 1$. Similarly, a sub_{if} transition is immediately followed by a ift , a iff or a ife transition. Therefore, $|\rho|_{\text{sub}_I} + |\rho|_{\text{sub}_{\text{if}}} \leq |\rho|_p + 1$.

Number of transitions: renaming steps. The analysis for $|\rho|_{\text{sub}_{\text{var}}}$ is subtler. A *variable crumble* is a crumble of the form (x, e) . The number of sub_{var} transitions is bounded by the number of variable crumbles out of bodies appearing in evaluation position along an execution $\rho : c_0 \rightarrow_{C_r}^* c$. These can be due to the following reasons:

- (1) *Static*: variable crumbles out of bodies in the initial state c_0 ;
- (2) *Dynamic*: variable crumbles obtained dynamically. In turn, these are divided into (see also the discussion after Prop. 5.7):
 - (a) *Copy*: variable crumbles occurring in the bodies of abstractions and if-then-else (and thus frozen) that become active because the construct is evaluated and the body exposed;
 - (b) *Creation*: variable crumbles that cannot be traced back to variable crumbles appearing in prefixes of the execution.

We now show that the crumbling translation does not produce any variable crumbles out of bodies, but one, if the original term is itself a variable. Therefore, the contribution of point 1 is at most 1. We need a measure, counting variable crumbles out of bodies. Note that a variable crumble (x, e) appearing in a crumble context C rather takes the form $[y \leftarrow x]e$, which is why the following measure counts the substitutions containing only a variable.

$$\begin{aligned} |b|_{\text{var}} &:= 0 && \text{if } b \text{ is not a variable} \\ |x|_{\text{var}} &:= 1 && |(b, e)|_{\text{var}} := |b|_{\text{var}} + |e|_{\text{var}} \\ |e|_{\text{var}} &:= 0 && |e[x \leftarrow b]|_{\text{var}} := |e|_{\text{var}} + |b|_{\text{var}}. \end{aligned}$$

PROPOSITION 5.7. *Let t be a term and v a value. Then:*

- (1) $|t|_{\text{var}} \leq 1$; and $|t|_{\text{var}} = 1$ if and only if t is a variable;
- (2) $|\bar{v}|_{\text{var}} \leq 1$; and $|\bar{v}|_{\text{var}} = 1$ if and only if v is a variable.

Proof
p. 21

Let us now discuss the variable crumbles of point 2.a (dynamic copy). By the bodies invariant (Lemma 5.5.3), these pairs appear in a body of the initial crumble. By the bodies property of the crumbling translation (Lemma 4.5.4), all these bodies are the translation of a term, and—by using Prop. 5.7 again—we obtain that each such body contributes at most with one variable crumble. Since each body is exposed by one \rightarrow_{β_v} or \rightarrow_{ift} or \rightarrow_{iff} transition, we have that the variable crumbles of point 2.a are bounded by $|\rho|_p$.

Last, we bound the number of variable crumbles at point 2.b (dynamic creation). There is only one rule that can create a new variable crumble (and exactly one), namely \rightarrow_{β_v} when the argument of the β -redex is a variable. For instance,

$$((\lambda x.(xx, \epsilon))y, [y \leftarrow \lambda z.z]) \rightarrow_{\beta_v} (xx, [x \leftarrow y][y \leftarrow \lambda z.z])$$

where the created variable crumble is $(y, [y \leftarrow \lambda z.z])$. Then the number of variable crumbles at point 2.b is bounded by the number of \rightarrow_{β_v} transitions, itself bounded by $|\rho|_p$.

The following lemma sums up the previous discussions

LEMMA 5.8. *Let $\rho : c_0 \rightarrow_{C_r}^* c$ be a Crumble GLAM execution.*

- (1) Linear number of non-renamings substitutions: $|\rho|_{\text{sub}_I} + |\rho|_{\text{sub}_{\text{if}}} \leq |\rho|_p + 1$.
- (2) Linear number of renamings: $|\rho|_{\text{sub}_{\text{var}}} \leq 2|\rho|_p + 1$.
- (3) Linear number of substitutions: $|\rho|_{\text{sub}_I} + |\rho|_{\text{sub}_{\text{var}}} + |\rho|_{\text{sub}_{\text{if}}} \leq 3|\rho|_p + 2$.

Cost of single transitions. Performing a single transition \rightarrow in the Crumble GLAM consists of four operations:

- (1) *Search*: locating the next redex;

- (2) *Unplugging*: splitting the crumble to be reduced into a crumble context C and the crumble c that is the redex at top level;
- (3) *Rewriting*: applying a rewriting rule to the crumble c , obtaining a new crumble d ;
- (4) *Plugging*: putting the new crumble back into the crumble context obtaining $C\langle d \rangle$.

The search for redexes is embedded into the definition of the rules, via the contextual closure. The technical definition of plugging and unplugging of crumbles into a crumble context is quite involved and, if implemented literally, is not constant time.

To ease the reasoning, in this section we assume that search and (un)plugging have negligible cost and show that the total cost of rewriting is bilinear. In Section 7 we sketch a slight variant of the Crumble GLAM, the Pointed Crumble GLAM, that adds a transition for searching redexes and removes the need for plugging and unplugging (details are in [?]). A further analysis of the Pointed Crumble GLAM shows that the total cost of search and (un)plugging is bilinear and thus negligible, justifying the results of this section.

Cost of single transitions: β_v transitions. We denote by $|t|$, $|c|$, $|e|$ and $|b|$ the size of terms, crumbles, environments and bites, respectively, defined as follows:

$$\begin{aligned}
 |x| &:= 1 & |tu| &:= |t| + |u| + 1 \\
 |\text{true}| &= |\text{false}| := 1 & |\lambda x.t| &:= |t| + 1 \\
 |\text{if } t \text{ then } u \text{ else } s| &:= |t| + |u| + |s| + 1 & |\text{err}| &:= 1 \\
 |e| &:= 0 & |e[x \leftarrow b]| &:= |e| + |b| \\
 |(b, e)| &:= |b| + |e|.
 \end{aligned}$$

The cost of each β_v transition (that needs to perform a copy of the crumble in the abstraction in order α -rename it) is bound by the size of the copied crumble. By the bodies invariant (Lemma 5.5.3) the abstraction is the α -renaming of one the abstractions already present in the initial crumble. Therefore the cost of a β_v transition is bound by the size of the initial crumble. The next lemma shows that the size of the initial crumble is linear in the size of the initial term translating to the crumble. Therefore, the cost of a β_v transition is linear by the size of the initial term.

LEMMA 5.9 (SIZE OF TRANSLATED TERMS). *Let t be a term and v a value. Then $|\underline{t}| \leq 5|t|$ and $|\overline{v}| \leq 5|v|$.*

Cost of single transitions: substitutions. The cost of sub_l , sub_{var} , and sub_{if} transitions depends on the choice of data structures for implementing the machine. Following the literature on global environment machines [5], we assume the global environment to be implemented as a store and variable occurrences to be implemented as pointers into the store, so that lookup in the environment can be performed in constant time on a Random Access Machine (RAM). As for the cost of actually performing the replacement of x with $e_v(x)$ in the sub_{var} , sub_l and sub_{if} rules, it can be done in constant time by copying the pointer to $e_v(x)$. This is possible because the actual copy, corresponding to α -renaming, is done in the β_v step. Thus, single substitution transitions have constant cost.

Cost of single transitions: conditionals and errors. It is immediate that—if one excludes plugging and unplugging—these transitions have constant cost.

TERMS	$t, u ::= \dots$	(as in $\lambda_{\text{Plot}}^{\text{if}}$, see Figure 1)
VALUES	$v ::= \dots$	(as in $\lambda_{\text{Plot}}^{\text{if}}$, see Figure 1)
FIREBALLS	$f ::= v \mid i$	
INERT TERMS	$i ::= x f \mid i f \mid \text{if } x \text{ then } t \text{ else } u \mid \text{if } i \text{ then } t \text{ else } u$	
RIGHT F-CONTEXT	$R ::= \langle \cdot \rangle \mid tR \mid Rf \mid \text{if } R \text{ then } u \text{ else } s$	
REDUCTION RULES AT TOP LEVEL		
$(\lambda x.t)i \mapsto_{\beta_i} t\{x \leftarrow i\}$	$\mapsto_{\beta_v}, \mapsto_{\text{ift}}, \mapsto_{\text{iff}}, \mapsto_{\text{ife}}, \mapsto_{@e}$	as in $\lambda_{\text{Plot}}^{\text{if}}$
CONTEXTUAL CLOSURE		
$R(t) \rightarrow_a R(u)$	if $t \mapsto_a u$	for $a \in \{\beta_v, \beta_i, \text{ift}, \text{iff}, \text{ife}, @e\}$
$\mapsto_{\beta_f} := \mapsto_{\beta_v} \cup \mapsto_{\beta_i}$	$\mapsto_{c\beta_f} := \bigcup_{a \in \{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e\}} \rightarrow_a$	

Figure 2: The conditional fireball calculus $\lambda_{\text{fire}}^{\text{if}}$.

Cost of executions. Summing up all the analyses in this section we obtain the following theorem.

THEOREM 5.10 (THE CRUMBLE GLAM IS BILINEAR UP TO SEARCH AND (UN)PLUGGING). *For any closed term t and any Crumble GLAM execution $\rho: \underline{t} \rightarrow_{\text{Cr}}^* c$, the cost of implementing ρ on a RAM is $O((|\rho|_p + 1) \cdot |t|)$ plus the cost of plugging and unplugging.*

OCaml implementation. In Section 7 we sketch the Pointed Crumble GLAM, a refinement of the Crumble GLAM making explicit the search for redexes and removing the need for (un)plugging, and having the *same complexity*: the cost for searching redexes and (un)plugging is negligible. More details and an implementation in OCaml of the Pointed Crumble GLAM can be found in the appendix of [?], together with the code that implements the crumbling translation. There we also discuss a parsimonious choice of data structures for the implementation of pointed environments.

6 THE OPEN CASE

6.1 The Fireball Calculus

In this section we recall the fireball calculus λ_{fire} , the simplest presentation of Open CbV, and extend it with conditionals. The extension is completely modular. For the issues of Plotkin's setting with respect to open terms and for alternative presentations of Open CbV, we refer the reader to Accattoli and Guerrieri [7, 9].

The fireball calculus was introduced without a name and studied first by Paolini and Ronchi Della Rocca [27], Ronchi Della Rocca and Paolini [29]. It has then been rediscovered by Grégoire and Leroy [19] to improve the implementation of Coq, and later by Accattoli and Sacerdoti Coen [10] to study cost models, where it was also named. We present it following Accattoli and Sacerdoti Coen [10], changing only inessential, cosmetic details.

The fireball calculus. The conditional fireball calculus $\lambda_{\text{fire}}^{\text{if}}$ is defined in Fig. 2. The conditional part is exactly as in the closed case. The idea is that the values of the Pif calculus are generalized to *fireballs*, by adding *inert terms*. Fireballs (noted f) and inert terms (noted i) are defined by mutual induction (in Fig. 2). For instance, x and $\lambda x.y$ are fireballs as values, while $y(\lambda x.x)$, xy , and $(z(\lambda x.x))(zz)(\lambda y.(zy))$ are fireballs as inert terms.

The main feature of inert terms is that they are open, normal, and that when plugged in a context they cannot create a redex, hence the name “inert”. Essentially, they are the *neutral terms* of Open

CbV. In Grégoire and Leroy's presentation [19], inert terms are called *accumulators* and fireballs are simply called values. Variables are, morally, both values and inert terms. In Accattoli and Sacerdoti Coen [10] they were considered as inert terms, while here, for minor technical reasons we prefer to consider them as values and not as inert terms—the change is inessential.

Evaluation rules. First, CbV β -reduction is replaced by *call-by-fireball* β -reduction \rightarrow_{β_f} : the β -rule can fire, *lighting* the argument, only if the argument is a fireball (*fireball* is a catchier version of *fire-able term*). We actually distinguish two sub-rules: the usual one that *lights* values, noted \rightarrow_{β_v} , and a new one that *lights* inert terms, noted \rightarrow_{β_i} (see Fig. 2). Second, we include all the rules about conditionals and errors, exactly as before, obtaining the evaluation relation $\rightarrow_{c\beta_f}$. Note that evaluation is *weak*: it does not reduce in abstraction nor if-then-else bodies.

We endow the calculus with the (deterministic) right-to-left evaluation strategy, defined via right f -contexts R —note the production Rf , forcing the right-to-left order. A more general calculus (without conditionals) is defined in Accattoli and Guerrieri [7], for which the right-to-left strategy is shown to be complete. We omit details about the rewriting theory of the fireball calculus because our focus here is on implementations.

Example 6.1. We have $(\lambda z.z(yz))(\lambda x.x) \rightarrow_{\beta_f} (\lambda x.x)(y(\lambda x.x)) \rightarrow_{\beta_f} y(\lambda x.x)$, where the final term $y(\lambda x.x)$ is a fireball (and β_f -normal).

Properties. As discussed in Sect. 5, Closed CbV enjoys harmony (Prop. 2.1). The fireball calculus $\lambda_{\text{fire}}^{\text{if}}$ satisfies an analogous property in the *open* setting by replacing abstractions with fireballs; we here further extend it to conditionals (Prop. 6.2.1 below). The key property of inert terms is summarized by Prop. 6.2.2: substitution of inert terms does not create or erase β_f -redexes, and hence can always be avoided. It plays a role in the design of the open abstract machine of the next section.

PROPOSITION 6.2 (PROPERTIES OF $\lambda_{\text{fire}}^{\text{if}}$). *Let t, u be terms.*

- (1) Open harmony: t is β_f -normal if and only if t is a fireball.
- (2) Inert substitutions and evaluation commute: *Let i be an inert term. Then $t \rightarrow_{c\beta_f} u$ if and only if $t\{x \leftarrow i\} \rightarrow_{c\beta_f} u\{x \leftarrow i\}$.*

6.2 The Open Crumble GLAM

Here we extend the Crumble GLAM defined in Sect. 5 to the case of open terms, implementing Open (Conditional) CbV, *i.e.* the conditional fireball calculus $\lambda_{\text{fire}}^{\text{if}}$: in this way we obtain the *Open Crumble GLAM*. The extension impacts on the core λ -calculus, while conditionals are essentially orthogonal to the issues of open terms.

Evaluated environments. First, we need to discuss the environments under which evaluation takes place. In the open case, v -crumbles and v -environments generalize to f -crumbles and f -environments, and are denoted as follows:

f -CRUMBLES: c_f f -ENVIRONMENTS: e_f

Recall that in the Crumble GLAM the already evaluated coda of the environment is made out only of practical values. Unfortunately, a syntactic characterization of f -environments (and f -crumbles) is more involved than the simple definition of v -environments.

In the Crumble GLAM, to check whether a bite b is in normal form with respect to a v -environment e_v , it suffices to check whether b is a practical value. In the open case, looking at the syntactic structure of the term is not enough: some applications are now normal, for example the bite yx is normal with respect to the environment $e := [x \leftarrow I]$, but not all of them are normal, for instance $(x\ y, [x \leftarrow I]) \rightarrow_{\text{sub}_I} (I\ y, [x \leftarrow I])$ as in the closed case (exact definitions are given below). Because of this additional complication, we are going to define f -environments directly in terms of their 'semantics', *i.e.* of their read-back to terms. Intuitively, fully evaluated f -environments should correspond to substitutions of fully evaluated terms in $\lambda_{\text{fire}}^{\text{if}}$. And since by harmony normal forms in $\lambda_{\text{fire}}^{\text{if}}$ are simply fireballs, it suffices to request that the read-back of every entry in a f -environment is a fireball.

Let us now define f -environments formally: e_f is a *f-environment* (resp. c_f is a *f-crumble*) if for any environment context E (resp. any crumble context C) and any crumble c such that $e_f = E\langle c \rangle$ (resp. $c_f = C\langle c \rangle$) the following two conditions hold:

- (1) *Read-back to fireballs*: $c \downarrow$ is a fireball, and
- (2) *Unchaining practical values*: if $c \downarrow$ is a practical value, then $c = (v, e)$ for some practical value v and some e .

The second requirement forbids v to be a variable and is crucial for capturing the correct behavior of the substitution rule $\rightarrow_{\text{sub}_{\text{var}}}$, which removes the malicious chains of substitutions (of variables for variables) discussed in Sect. 5.

Transitions. The transitions of the Open Crumble GLAM:

$$((\lambda x.c)v, e_f) \mapsto_{\beta_f} (c @ [x \leftarrow v])^\alpha @ e_f$$

$$(\text{if true then } c \text{ else } d, e_f) \mapsto_{\text{ift}} c @ e_f$$

$$(\text{if false then } c \text{ else } d, e_f) \mapsto_{\text{iff}} d @ e_f$$

$$(\text{if } v \text{ then } c \text{ else } d, e_f) \mapsto_{\text{ife}} (\text{err}, e_f) \quad (1)$$

$$(v w, e_f) \mapsto_{@e} (\text{err}, e_f) \quad (2)$$

$$(x, e_f) \mapsto_{\text{sub}_{\text{var}}} (e_f(x), e_f) \quad (3)$$

$$(x v, e_f) \mapsto_{\text{sub}_I} (e_f(x) v, e_f) \quad (3)$$

$$(\text{if } x \text{ then } c \text{ else } d, e_f) \mapsto_{\text{sub}_{\text{if}}} (\text{if } e_f(x) \text{ then } c \text{ else } d, e_f) \quad (3)$$

$$(1) \text{ if } v = \lambda x.e \text{ or } v = \text{err}$$

$$(2) \text{ if } v \in \{\text{true}, \text{false}, \text{err}\}$$

$$(3) \text{ if } x \in \text{dom}(e_f)$$

Top level transitions are then closed by crumble contexts by setting $C\langle c \rangle \rightarrow_a C\langle d \rangle$ if $c \mapsto_a d$ for $a \in \{\beta_v, \text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}, \text{ift}, \text{iff}, \text{ife}, @e\}$. The transition relation \rightarrow_{oCr} of the Open Crumble GLAM is defined as the union of all these rules. A *principal transition* of the Open Crumble GLAM is a transition \rightarrow_a for any rule $a \in \{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e\}$.

There are only two differences with the transitions of the Crumble GLAM. First, \rightarrow_{β_v} is now noted \mapsto_{β_f} and yet it is identical to the one in the closed case (the comments about α -renaming given in Sect. 5 still hold). This is because there is a subtle difference: the argument of the β -redex may be a variable (which is a value) substituted by a inert term in the environment, thus becoming a \rightarrow_{β_i} step (and not a \rightarrow_{β_v} step) when read-back in $\lambda_{\text{fire}}^{\text{if}}$. Second, there is a slightly different side condition for the substitution transitions: it requires not only that a variable is defined in e_f (like in the closed

case), but also that the corresponding term in the environment is a practical value (and not an inert term nor a variable).

Note that the substitution transitions substitute values only. The environment e_f may contain also bites that are variables or applications, but these bites are not substituted: this choice is justified by the property of $\lambda_{\text{fire}}^{\text{if}}$ stated in Prop. 6.2.2. Besides, avoiding the substitution of inert terms is a prerequisite for efficiency of the machine, that would otherwise be subjected to an exponential overhead due to *size explosion*, see for example Accattoli and Guerrieri [8], Accattoli and Sacerdoti Coen [10].

The harmony between evaluation rules and the syntactic definition of normal forms is witnessed by the following property.

PROPOSITION 6.3 (HARMONY FOR THE OPEN CRUMBLE GLAM). *A crumble c is oCr-normal if and only if it is a f -crumble.*

Example 6.4. Recall that $\bar{\delta} = (\lambda x.xx, \epsilon)$. In Example 5.2 we noted that the (open) crumble $\delta\delta(xx)$ was stuck in the Crumble GLAM. Now instead it correctly reduces, never reaching a normal form:

$$\begin{aligned} \delta\delta(xx) &= (zw, [z \leftarrow \bar{\delta}][w \leftarrow xx]) \\ &\rightarrow_{\beta_v} (zw, [z \leftarrow yy][y \leftarrow \bar{\delta}][w \leftarrow xx]) \\ &\rightarrow_{\text{sub}_l} (zw, [z \leftarrow \bar{\delta}y][y \leftarrow \bar{\delta}][w \leftarrow xx]) \rightarrow \dots \end{aligned}$$

Implementation Theorem. The proof of the implementation theorem for the Open Crumble GLAM follows the same structure as for the Crumble GLAM in Subsect. 5.2, relying on similar but subtler invariants that can be found in the appendix of [?].

THEOREM 6.5 (IMPLEMENTATION). *Let c be a crumble that is reachable by the Open Crumble GLAM.*

- (1) Initialization: $t_{\downarrow} = t$
- (2) Principal projection: *if $c \rightarrow_a d$ then $c_{\downarrow} \rightarrow_a d_{\downarrow}$ for $a \in \{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e\}$.*
- (3) Overhead transparency: *if $c \rightarrow_a d$ then $c_{\downarrow} = d_{\downarrow}$ for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$.*
- (4) Determinism: *the transition \rightarrow_{oCr} is deterministic.*
- (5) Halt: *if c is $\rightarrow_{\text{oCr-normal}}$ then c_{\downarrow} is $\rightarrow_{\text{c}\beta_f}$ -normal.*
- (6) Overhead termination: *\rightarrow_a terminates, for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$.*

Therefore, the Open Crumble GLAM, the right-to-left conditional fire-ball evaluation $\rightarrow_{\text{c}\beta_f}$ and the read-back $(\cdot)_{\downarrow}$ form an implementation system.

Complexity. The complexity analysis is identical to the one in Subsect. 5.3. Indeed, once the search for the next redex and (un)plugging are neglected, the two machines only differ by the additional $O(1)$ side condition for the substitution transitions.

THEOREM 6.6 (THE OPEN CRUMBLE GLAM IS BILINEAR UP TO SEARCH AND (UN)PLUGGING). *For any term t and any Open Crumble GLAM execution $d: t \rightarrow_{\text{oCr}}^* c$, the cost of implementing ρ on a RAM is $O((|\rho|_p + 1) \cdot |t|)$ plus the cost of plugging and unplugging.*

OCaml implementation. Following the same pattern of the closed case, in Section 7 we introduce a machine making explicit the search for redexes and removing the need of (un)plugging, so as to show that their cost is negligible. The OCaml code implementing this further machine is in the appendix of [?], along with a detailed discussion of the adopted data structures. The code for the open and

closed machines is identical but for five lines: three implement the additional check for practical values in the substitution transitions, the others consider also inert terms in the search transition.

7 THE (OPEN) POINTED CRUMBLE GLAM

In the abstract machines considered so far the *search for the next redex* is implicit in definition of the evaluation rules, as it corresponds to applying rules to crumbles with already evaluated environments and to factoring out the crumble context.

Here we sketch how to make the search explicit by introducing a variant of the Crumble GLAM called *Pointed Crumble GLAM*. The new search transitions have constant cost and the machine is bilinear—the detailed study is in the appendix of [?].

Pointed crumbles and pointed environments. The key idea behind the Pointed Crumble GLAM is to avoid (un)plugging in the rules by letting them act on *pointed crumbles*, i.e. on crumbles where a *pointer* marks explicitly the dividing point between the evaluated coda and the crumbled term of the currently active crumble. A pointed crumble $(b, e[x \leftarrow b'] \mid e_v)$ represents the crumble $C((b', e_v))$, where $C = (b, e[x \leftarrow (\cdot)])$ is the crumble context, (b', e_v) is the active crumble, and e_v is the evaluated coda. If (b', e_v) is a Crumble GLAM a -redex (for rule $a \in \{\beta_v, \text{sub}_{\text{var}}, \text{sub}_l\}$), the Pointed Crumble GLAM shall reduce according to the corresponding a -transition that also takes care of setting (in $O(1)$) the pointer to the rightmost unevaluated crumble. Otherwise, by harmony (Prop. 5.3), b' must be a crumbled value v and therefore the pointer is moved (in $O(1)$) one step to the left, looking for the next redex, via the search transition $(b, e[x \leftarrow v] \mid e_v) \rightarrow_{\text{sea}} (b, e \mid [x \leftarrow v]e_v)$.

Unfortunately, there is an annoying technical issue. Not all pointed configurations are of the form $(b, [x \leftarrow b'] \mid e_v)$: the configurations $(b, \mid e_v)$ must be also taken into account and reduced if b is not a crumbled value. However, there is no simple way to describe transitions that act uniformly on configurations $(b, \mid e_v)$ and $(b, e[x \leftarrow b'] \mid e_v)$ without duplicating the rules or without reintroducing a notion of contextual closure. To solve the issue, we abandon pointed crumbles and adopt pointed environments instead.

A *pointed environment* $([x \leftarrow b]e \mid e_v)$ is just a representation of a pointed crumble $(b, e \mid e_v)$. The leftmost variable x in a pointed environment can be understood as the name given to the machine output. It plays a role similar to the outermost λ -abstraction introduced by CPS translations, that binds the continuation that is fed with the output of the evaluation. In particular, a normal pointed environment $(\mid [x \leftarrow v]e_v)$ represents the normal crumble (v, e_v) .

Formal definition of pointed environments and transitions. Pointed environments are defined as $e_{\downarrow} := e \mid e'$ for any non-pointed environments e and e' where at least one among e or e' is non-empty. The translation $\iota(\cdot)$ embeds crumbles into pointed environments: $\iota(b, e) := [x \leftarrow b]e \mid \epsilon$, where x is any variable fresh in b and e .

The transition rules of the Pointed Crumble GLAM are:

$$\begin{aligned} e[x \leftarrow (\lambda y.c)v] \mid e_v &\rightarrow_{\beta_v} e[x \leftarrow b]e'[z \leftarrow v] \mid e_v & (i) \\ e[y \leftarrow x] \mid e_v &\rightarrow_{\text{sub}_{\text{var}}} e[y \leftarrow e_v(x)] \mid e_v & (ii) \\ e[y \leftarrow xv] \mid e_v &\rightarrow_{\text{sub}_l} e[y \leftarrow e_v(x)v] \mid e_v & (ii) \\ e[x \leftarrow b] \mid e_v &\rightarrow_{\text{sea}} e \mid [x \leftarrow b]e_v & (iii) \end{aligned}$$

where (we omitted the rules for conditionals, for the sake of brevity)

Proof
p. 27

Proof
p. 29

- i. $\lambda z.(b, e') := (\lambda y.c)^\alpha$ with $(e[x \leftarrow b]e'[z \leftarrow v] \vdash e_v)$ well-named;
- ii. if $x \in \text{dom}(e_v)$;
- iii. if none of the other rules is applicable, i.e. when b is an abstraction or when b is x or xv but x is not defined in e_v .

In the appendix of [?], we prove that the Pointed Crumble GLAM simulates the Crumble GLAM following exactly the same schema already used in the paper, namely they form an implementation system. We also provide the complexity analysis, and smoothly lift everything to the open case, by studying the *Open Pointed Crumble GLAM* (which simulates the Open Crumble GLAM and whose transition function is noted $\rightarrow_{\text{poCr}}$). We obtain the following result, that sums up the study in this paper (Point 2 is a corollary of Point 1).

THEOREM 7.1 (THE OPEN (POINTED) CRUMBLE GLAM IS BILINEAR).
Let t be a term.

- (1) for any Open Pointed Crumble GLAM execution $\rho: \iota(t) \rightarrow_{\text{poCr}}^* e_1$, the cost of implementing ρ on a RAM is $O((|\rho|_p + 1) \cdot |t|)$.
- (2) for any Open Crumble GLAM execution $\rho: \underline{t} \rightarrow_{\text{oCr}}^* c$, the cost of implementing ρ on a RAM is $O((|\rho|_p + 1) \cdot |t|)$.

8 EXTENSIONS

Left-to-right CbV. The (right-to-left) Crumble GLAM can also implement a left-to-right strategy for the Pif calculus. The only change concerns the crumbling transformation, that on applications has to put the environment coming from the (transformation of the) left subterm on the right of the one coming from the right subterm.

Call-by-need. The crumbling technique applies also to call-by-need machines. There are however a few differences. First, the machine does no longer explore sequentially the environment from right-to-left, it rather starts on the left and then jumps back and forth, by need. Then the definition of evaluation contexts is trickier, especially in the open case.

Strong CbV. Simply designing an abstract machine for strong reduction is relatively easy. However the easy machines are not *bilinear*, and not even polynomial.

The needed optimizations to make them reasonable (i.e. polynomial or bilinear) are clear, they are the same at work in the open case (or in the call-by-name case):

- (1) substitute only abstraction and not inert terms, and
- (2) do not substitute abstractions on variable occurrences that are not applied.

These principles however have different consequences in different settings. In particular, (2) implies that some abstraction are kept shared forever, and a *strong* CbV approach has to evaluate them (while the open setting does not) and only once, thus it has to evaluate them while they are shared, adding a call-by-need flavor.

There are two difficulties. First, the specification of the search for redexes, that becomes involved and requires many machine transitions—the crumbling technique is meant to help here. Second, the proof of correctness of the machine.

All proofs of correctness in the literature (including those in this paper) are simulations up to sharing based on a *bijection* of β -redexes (or principal steps) between the abstract machine and the λ -calculus strategy (one half of the bijection is the principal

projection property of implementation systems in Section 5.2, the other half is implied by the other properties).

The evaluation under *shared* abstraction required by CbV strong evaluation breaks the usual bijection of β -redexes (as one β -transition of the machine is mapped to *many* β -steps on the calculus, and not necessarily those of a standard strategy), thus forbidding to employ the standard technique for proving correctness.

The new proof technique for correctness for reasonable strong CbV and the intricacies of the search for redexes in the strong case, do deserve to be studied carefully, and are thus left to future work.

9 COMMENTS ON RELATED WORKS

Here we discuss Kennedy's potential slowdown and provide a counterexample to the scalability of the CPS translation to open terms.

Kennedy. Kennedy [21] compares three different calculi: a monadic calculus, which has ES, a calculus of administrative normal forms (ANFs) and the image of a CPS transformation. In the monadic calculus β_v -redexes can be hidden by ES which need to be commuted to reveal the β_v -redex. Kennedy shows an example (see Fig. 3) where the number of commutations is not bounded linearly by the number of β_v -steps and blames the inefficiency of his compiler on that. In his example, the number of commutations is *quadratic* in the number of β_v -steps, since the i^{th} β_v -step is immediately followed by i commutation steps.

ANFs are just canonical shapes of monadic terms where the top-most term and the body of each abstraction is a crumble, i.e. a term together with a list of ES that map variables to terms (instead of crumbles). Kennedy rightly observes that ANFs are not preserved by standard β_v -reduction, and thus, after each β_v -step, some commutative steps are required to reach the ANF shape. Kennedy too hastily concludes that the quadratic blowup also affects the ANF calculus, since its quadratic example stands in the ANF fragment.

However, Kennedy misses the fact that the ES in ANFs form a list and that the commutations steps altogether just implement the append of two lists. Since append can be implemented in constant time, the complexity of evaluation in the ANF calculus is just linear (and not quadratic) in the number of β_v -steps. This is the same complexity we achieved for the Crumble and Open Crumble GLAM.

Danvy and Filinski. In Danvy and Filinski [16] the CPS transformation is shown to scale up to open terms (their Theorem 2). On open terms, however, they consider Plotkin's CbV operational semantics λ_{Plot} , which is *not adequate* (it is adequate only for closed terms, see Accattoli and Guerrieri [7, 9] and ?). When one considers one of the equivalent adequate CbV semantics in Accattoli and Guerrieri [7, 9] for the open case, for instance the fireball calculus λ_{fire} , then the properties of the CPS no longer hold, in particular it does not commute with evaluation, as the following example shows. Take the following open term $t := (\lambda x. \lambda y. y)(zz)v$, where v is a value, say a distinguished variable. In λ_{Plot} the term t is β_v -normal, but in λ_{fire} we have:

$$t := (\lambda x. \lambda y. y)(zz)v \rightarrow_{\beta_f} (\lambda y. y)v \rightarrow_{\beta_f} v$$

Now, consider the CPS translation $\text{cps}(t)$ of t , according to the definition in Danvy and Filinski [16]. We use λ for standard ("dynamic", in Danvy's terminology) abstraction, and Λ and $@$ for "static" abstraction and "static" prefix application, respectively. If a

$$\begin{aligned}
t &:= (z_1 x_0)[z_1 \leftarrow \lambda x_1. b y_1[y_1 \leftarrow z_2 x_1]][z_2 \leftarrow \lambda x_2. b y_2[y_2 \leftarrow z_3 x_2]] \dots [z_n \leftarrow \lambda x_n. b y_n[y_n \leftarrow b x_n]] \\
&\xrightarrow{(\rightarrow_{\beta_v} \rightarrow_{\text{let}}) (\rightarrow_{\beta_v} \rightarrow_{\text{let}}^2) \dots (\rightarrow_{\beta_v} \rightarrow_{\text{let}}^i)} \\
&\quad (z_1 x_0)[z_1 \leftarrow \lambda x_1. b y_1[y_1 \leftarrow z_2 x_1]][z_2 \leftarrow \lambda x_2. b y_2[y_2 \leftarrow z_3 x_2]] \dots \\
&\quad \dots [z_{n-i-1} \leftarrow \lambda x_{n-i-1}. b y_{n-i-1}[y_{n-i-1} \leftarrow z_{n-i} x_{n-i-1}]] [z_{n-i} \leftarrow \lambda x_{n-i}. b y_{n-i}[y_{n-i} \leftarrow b y_{n-i+1}][y_{n-i+1} \leftarrow b y_{n-i+2}] \dots [y_n \leftarrow b x_{n-i}]] \\
&\xrightarrow{\beta_v} \\
&\quad (z_1 x_0)[z_1 \leftarrow \lambda x_1. b y_1[y_1 \leftarrow z_2 x_1]][z_2 \leftarrow \lambda x_2. b y_2[y_2 \leftarrow z_3 x_2]] \dots \\
&\quad \dots [z_{n-i-1} \leftarrow \lambda x_{n-i-1}. b y_{n-i-1}[y_{n-i-1} \leftarrow b y_{n-i}[y_{n-i} \leftarrow b y_{n-i+1}][y_{n-i+1} \leftarrow b y_{n-i+2}] \dots [y_n \leftarrow b x_{n-i-1}]]] \\
&\xrightarrow{\text{let}^{i+1}} \\
&\quad (z_1 x_0)[z_1 \leftarrow \lambda x_1. b y_1[y_1 \leftarrow z_2 x_1]][z_2 \leftarrow \lambda x_2. b y_2[y_2 \leftarrow z_3 x_2]] \dots \\
&\quad \dots [z_{n-i-1} \leftarrow \lambda x_{n-i-1}. b y_{n-i-1}[y_{n-i-1} \leftarrow b y_{n-i}][y_{n-i} \leftarrow b y_{n-i+1}] \dots [y_n \leftarrow b x_{n-i-1}]]
\end{aligned}$$

Figure 3: Kennedy’s example of evaluation in the monadic calculus where the number of commutation steps is quadratic in the number of β_v -steps (\rightarrow^i stands for the composition of $i \rightarrow$ -steps). The i^{th} β_v -step (which can reduce under abstractions) is immediately followed by i commutation steps \rightarrow_{let} that just append two lists of substitutions moving one substitution at a time. Thus, to reach a normal form one needs $n \beta_v$ -steps and $n(n+1)/2$ let-steps. In the Crumble and Open Crumble GLAM instead, the commutation steps are integrated in the β_v -rule simply by appending the two lists in constant time.

generalized version of Theorem 2 in Danvy and Filinski [16] held in the open case, one would expect that $@(\text{cps}(t))I$ (where $I := \lambda z.z$) evaluates to v , as v is a value. But, even using an unrestricted β -reduction that goes under abstraction as evaluation, we obtain (we reduce all static redexes first, followed by all dynamic redexes):

$$\begin{aligned}
& @(\text{cps}(t))I \\
&= (\lambda k. @(\lambda x. @(\lambda y. @y(\lambda w. \lambda a. @(\lambda b. @b(\lambda c. \lambda d. @(\lambda e. @ec)(\lambda e. de))) \\
&\quad (\lambda b. ab))) @(\lambda j. @(\lambda a. @az)(\lambda a. @(\lambda b. @bz)(\lambda b. (ab)(\lambda c. @jc)))) \\
&\quad (\lambda w. (yw)(\lambda a. @xa))) @(\lambda x. @(\lambda y. @yv)(\lambda y. (xy)(\lambda w. @Kw))))I \\
&\xrightarrow{\beta^*} (zz)(\lambda x. ((\lambda y. \lambda w. w(\lambda a. \lambda b. ba))x)(\lambda y. yv(\lambda w. Iw))) \\
&\xrightarrow{\beta^*} (zz)(\lambda x. v)
\end{aligned}$$

where $(zz)(\lambda x. v)$ is not even β -equivalent to v . The CPS translation—like Plotkin’s calculus—gets stuck trying to evaluate zz , whereas the term reduces to v in the fireball calculus.

Summing up, we are not claiming that Theorem 2 in Danvy and Filinski [16] is false, but just that it does not mean that their CPS transformation scales up to open terms: to prove scalability, one should use an adequate CbV evaluation for open terms (such as the one of the fireball calculus), instead of Plotkin’s one. Our counterexample shows that Danvy’s and Filinski’s CPS does not scale up to open terms with an adequate CbV operational semantics for them.

This problem affects also other CPS translations, such as the ones defined by Plotkin [28] or by Lassen [23]. Likely, this is the reason why Lassen [23] states his Theorem 4.6 (the analogous of Theorem 2 in Danvy and Filinski [16]) only for closed terms.

10 CONCLUSIONS

This paper studies abstract machines working on crumbled forms with respect to design, efficiency, scalability, and implementations, putting emphasis on the role played by environments and providing a detailed technical development. In particular, we study the crumble setting on top of global environments—in future work we would like to explore the more technical case of local environments.

At the level of design, switching to crumbled forms removes the need for machine data structures such as the applicative stack or the dump, as they are encoded in crumbled environments.

At the level of efficiency, the evaluation of crumbled forms does not require any overhead: crumble abstract machines are linear in the number of steps of the calculus and in the size of the initial term, exactly as ordinary abstract machines with global environments.

At the level of scalability, everything—including the complexity—smoothly scales up from the closed case, relevant for programming languages, to the more delicate case of open terms, needed to implement proof assistants. As shown in Section 9, CPS translations do not smoothly scale up to the open case (contrary to what claimed by Danvy and Filinski [16]), so that our work shows an advantage of the crumbling transformation in this setting.

At the level of implementations, we stress the different operations on crumbled environments (sequential access and concatenation) and provide a concrete implementation, which does not suffer from the potential slowdown of crumbled forms pointed out by Kennedy [21] (see Section 9).

In future work we plan to apply our results to the design of abstract machines for strong call-by-value and call-by-need evaluation. Preliminary results suggest that the simplification to the code noticed in the open case is preserved and even amplified in the harder case of strong evaluation.

Acknowledgments. This work has been partially funded by the ANR JCJC grant COCA HOLA (ANR-16-CE40-004-01) and by the EPSRC grant EP/R029121/1 “Typed Lambda-Calculi with Sharing and Unsharing”.

REFERENCES

- [1] Beniamino Accattoli. 2012. An Abstract Factorization Theorem for Explicit Substitutions. In *23rd International Conference on Rewriting Techniques and Applications, RTA 2012 (LIPIcs)*, Vol. 15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 6–21. <https://doi.org/10.4230/LIPIcs.RTA.2012.6>
- [2] Beniamino Accattoli. 2016. The Useful MAM, a Reasonable Implementation of the Strong λ -Calculus. In *Logic, Language, Information, and Computation -*

- 23rd International Workshop, WoLLIC 2016 (Lecture Notes in Computer Science), Vol. 9803. Springer, 1–21. https://doi.org/10.1007/978-3-662-52921-8_1
- [3] Beniamino Accattoli, Pablo Barenbaum, and Damiano Mazza. 2014. Distilling abstract machines. In *19th ACM SIGPLAN International Conference on Functional Programming, ICFP 2014*. ACM, 363–376. <https://doi.org/10.1145/2628136.2628154>
 - [4] Beniamino Accattoli, Pablo Barenbaum, and Damiano Mazza. 2015. A Strong Distillery. In *Programming Languages and Systems - 13th Asian Symposium, APLAS 2015 (Lecture Notes in Computer Science)*, Vol. 9458. Springer, 231–250. https://doi.org/10.1007/978-3-319-26529-2_13
 - [5] Beniamino Accattoli and Bruno Barras. 2017. Environments and the Complexity of Abstract Machines. In *19th International Symposium on Principles and Practice of Declarative Programming, PPDP 2017*. ACM, 4–16. <https://doi.org/10.1145/3131851.3131855>
 - [6] Beniamino Accattoli and Ugo Dal Lago. 2016. (Leftmost-Outermost) Beta Reduction is Invariant, Indeed. *Logical Methods in Computer Science* 12, 1 (2016). [https://doi.org/10.2168/LMCS-12\(1:4\)2016](https://doi.org/10.2168/LMCS-12(1:4)2016)
 - [7] Beniamino Accattoli and Giulio Guerrieri. 2016. Open Call-by-Value. In *Programming Languages and Systems - 14th Asian Symposium, APLAS 2016 (Lecture Notes in Computer Science)*, Vol. 10017. Springer, 206–226. https://doi.org/10.1007/978-3-319-47958-3_12
 - [8] Beniamino Accattoli and Giulio Guerrieri. 2017. Implementing Open Call-by-Value. In *Fundamentals of Software Engineering - 7th International Conference, FSEN 2017 (Lecture Notes in Computer Science)*, Vol. 10522. Springer, 1–19. https://doi.org/10.1007/978-3-319-68972-2_1
 - [9] Beniamino Accattoli and Giulio Guerrieri. 2018. Types of Fireballs. In *Programming Languages and Systems - 16th Asian Symposium, APLAS 2018 (Lecture Notes in Computer Science)*, Vol. 11275. Springer, 45–66. https://doi.org/10.1007/978-3-030-02768-1_3
 - [10] Beniamino Accattoli and Claudio Sacerdoti Coen. 2015. On the Relative Usefulness of Fireballs. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015*. IEEE Computer Society, 141–155. <https://doi.org/10.1109/LICS.2015.23>
 - [11] Beniamino Accattoli and Claudio Sacerdoti Coen. 2017. On the value of variables. *Information and Computation* 255 (2017), 224–242. <https://doi.org/10.1016/j.ic.2017.01.003>
 - [12] Zena M. Ariola, Matthias Felleisen, John Maraist, Martin Odersky, and Philip Wadler. 1995. The Call-by-Need Lambda Calculus. In *22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL'95*. ACM Press, 233–246. <https://doi.org/10.1145/199448.199507>
 - [13] Pierre Crégut. 1990. An Abstract Machine for Lambda-Terms Normalization. In *LISP and Functional Programming*. 333–340. <https://doi.org/10.1145/91556.91681>
 - [14] Olivier Danvy. 1994. Back to Direct Style. *Science of Computer Programming* 22, 3 (1994), 183–195. [https://doi.org/10.1016/0167-6423\(94\)00003-4](https://doi.org/10.1016/0167-6423(94)00003-4)
 - [15] Olivier Danvy. 2003. A New One-Pass Transformation into Monadic Normal Form. In *Compiler Construction, 12th International Conference, CC 2003 (Lecture Notes in Computer Science)*, Vol. 2622. Springer, 77–89. https://doi.org/10.1007/3-540-36579-6_6
 - [16] Olivier Danvy and Andrzej Filinski. 1992. Representing Control: A Study of the CPS Transformation. *Mathematical Structures in Computer Science* 2, 4 (1992), 361–391. <https://doi.org/10.1017/S0960129500001535>
 - [17] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. 1993. The essence of compiling with continuations (with retrospective). In *20 Years of the ACM SIGPLAN Conference on Programming Language Design and Implementation 1979-1999, A Selection, PLDI 1993*. ACM, 502–514. <https://doi.org/10.1145/989393.989443>
 - [18] Álvaro García-Pérez, Pablo Nogueira, and Juan José Moreno-Navarro. 2013. Deriving the full-reducing Krivine machine from the small-step operational semantics of normal order. In *15th International Symposium on Principles and Practice of Declarative Programming, PPDP'13*. ACM, 85–96. <https://doi.org/10.1145/2505879.2505887>
 - [19] Benjamin Grégoire and Xavier Leroy. 2002. A compiled implementation of strong reduction. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming, ICFP '02*. ACM, 235–246. <https://doi.org/10.1145/581478.581501>
 - [20] John Hatcliff and Olivier Danvy. 1994. A Generic Account of Continuation-Passing Styles. In *21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 1994*. ACM Press, 458–471. <https://doi.org/10.1145/174675.178053>
 - [21] Andrew Kennedy. 2007. Compiling with continuations, continued. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming, ICFP 2007*. ACM, 177–190. <https://doi.org/10.1145/1291151.1291179>
 - [22] Arne Kutzner and Manfred Schmidt-Schauß. 1998. A Non-Deterministic Call-by-Need Lambda Calculus. In *Third ACM SIGPLAN International Conference on Functional Programming, ICFP 1998*. ACM, 324–335. <https://doi.org/10.1145/289423.289462>
 - [23] Søren B. Lassen. 2005. Eager Normal Form Bisimulation. In *20th IEEE Symposium on Logic in Computer Science, LICS 2005*. IEEE Computer Society, 345–354. <https://doi.org/10.1109/LICS.2005.15>
 - [24] John Launchbury. 1993. A Natural Semantics for Lazy Evaluation. In *Twentieth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 1993*. ACM Press, 144–154. <https://doi.org/10.1145/158511.158618>
 - [25] John Maraist, Martin Odersky, and Philip Wadler. 1998. The Call-by-Need Lambda Calculus. *Journal of Functional Programming* 8, 3 (1998), 275–317.
 - [26] Eugenio Moggi. 1991. Notions of Computation and Monads. *Information and Computation* 93, 1 (1991), 55–92. [https://doi.org/10.1016/0890-5401\(91\)90052-4](https://doi.org/10.1016/0890-5401(91)90052-4)
 - [27] Luca Paolini and Simona Ronchi Della Rocca. 1999. Call-by-value Solvability. *ITA* 33, 6 (1999), 507–534. <https://doi.org/10.1051/ita:1999130>
 - [28] Gordon D. Plotkin. 1975. Call-by-Name, Call-by-Value and the lambda-Calculus. *Theoretical Computer Science* 1, 2 (1975), 125–159. [https://doi.org/10.1016/0304-3975\(75\)90017-1](https://doi.org/10.1016/0304-3975(75)90017-1)
 - [29] Simona Ronchi Della Rocca and Luca Paolini. 2004. *The Parametric λ -Calculus – A Metamodel for Computation*. Springer. <https://doi.org/10.1007/978-3-662-10394-4>
 - [30] Amr Sabry and Matthias Felleisen. 1993. Reasoning about Programs in Continuation-Passing Style. *Lisp and Symbolic Computation* 6, 3-4 (1993), 289–360.
 - [31] Peter Sestoft. 1997. Deriving a Lazy Abstract Machine. *Journal of Functional Programming* 7, 3 (1997), 231–264.
 - [32] Christopher P. Wadsworth. 1971. *Semantics and pragmatics of the lambda-calculus*. PhD Thesis. Oxford. Chapter 4.

TECHNICAL APPENDIX

A PROOFS OF SECT. 2 (PIF CALCULUS)

PROOF OF (PROP. 2.1). Let t be a closed term. t is \rightarrow_{pif} -normal if and only if t is a value.

PROOF.

(\Rightarrow) Proof by induction on the structure of t .

- Case t value: trivial.
- Case $t = us$ for some terms u and s : we show that this case is not possible by deriving a contradiction. Since t is \rightarrow_{pif} -normal, then s is \rightarrow_{pif} -normal, and hence by *i.h.* s is a value. Since t is \rightarrow_{pif} -normal and s is a value, then also u must be \rightarrow_{pif} -normal, and hence a value. We proceed by cases on u , showing that no case is possible. u cannot be a variable, because t is closed; it cannot be an abstraction (because otherwise the rule \rightarrow_{β_v} may be applied, contradicting the hypothesis that t is \rightarrow_{pif} -normal); it cannot be a boolean or err (because otherwise the rule $\rightarrow_{@e}$ may be applied).
- Case $t = \text{if } u \text{ then } s \text{ else } r$ for some terms u, s, r : we show that this case is not possible by deriving a contradiction. Since t is \rightarrow_{pif} -normal, then u is \rightarrow_{pif} -normal, and hence by *i.h.* u is a value. We proceed by cases on u , showing that no case is possible. u cannot be a variable, because t is closed; it cannot be an abstraction or err (because otherwise the rule \rightarrow_{ife} may be applied, contradicting the hypothesis that t is \rightarrow_{pif} -normal); it cannot be a boolean (because otherwise one of the rules \rightarrow_{ift} or \rightarrow_{iff} may be applied).

(\Leftarrow) By hypothesis, t is a value. We proceed by cases on t . t cannot be a variable, because by hypothesis t is closed. If t is an abstraction, then it is \rightarrow_{pif} -normal since \rightarrow_{pif} does not reduce under λ 's. Otherwise t is either a boolean or err, which clearly are \rightarrow_{pif} -normal. \square

LEMMA A.1 (COMPOSITION OF RIGHT V-CONTEXTS). Let R and R' be right v-contexts. Then their composition $R\langle R' \rangle$ is a right v-context.

PROOF. By induction on the right v-context R . Cases:

- Hole, i.e. $R := \langle \cdot \rangle$: then, $R\langle R' \rangle = R'$ is a right v-context by hypothesis.
- Right, i.e. $R := tR''$: then, $R\langle R' \rangle = tR''\langle R' \rangle$ is a right v-context because $R''\langle R' \rangle$ is a right v-context by *i.h.*
- Left, i.e. $R := R''v$: then, $R\langle R' \rangle = R''\langle R' \rangle v$ is a right v-context because $R''\langle R' \rangle$ is a right v-context by *i.h.*
- if-then-else, i.e. $R := (\text{if } R'' \text{ then } t \text{ else } u)$: then, $R\langle R' \rangle = (\text{if } R''\langle R' \rangle \text{ then } t \text{ else } u)$ is a right v-context because $R''\langle R' \rangle$ is a right v-context by *i.h.* \square

B PROOFS OF SECT. 4 (PRELIMINARIES)

LEMMA B.1. Let $e := [x_1 \leftarrow b_1] \dots [x_n \leftarrow b_n]$ a well-named environment. If $\text{fv}(t_i) \cap \{x_1, \dots, x_i\} = \emptyset$ for all $1 \leq i \leq n$, then $\text{fv}(e) = \bigcup_{i=1}^n \text{fv}(b_i) \setminus \text{dom}(e)$ and $\text{fv}((b, e)) = (\text{fv}(b) \cup \text{fv}(e)) \setminus \text{dom}(e)$.

Definition B.2 (Disjointedness). Let P and P' be two crumbles or environments: P and P' are *disjoint* (noted $P \# P'$) if $\text{fv}(P) \cap \text{dom}(P') = \emptyset$.

Definition B.3 (Composition of crumble contexts). It is also possible to plug a crumble context C' into another crumble context C , as follows:

$$C\langle C' \rangle := \begin{cases} C & \text{if } C' = \langle \cdot \rangle \\ C' & \text{if } C = \langle \cdot \rangle \\ (b, e[x \leftarrow b']e'[y \leftarrow \langle \cdot \rangle]) & \text{if } C = (b, e[x \leftarrow \langle \cdot \rangle]) \text{ and } C' = (b', e'[y \leftarrow \langle \cdot \rangle]) \end{cases}$$

The *appending* of an environment context E to a crumble (b, e') is defined as $(b, e') @ E := (b, e'E)$.

The notions of $\text{fv}(\cdot)$ and $\text{dom}(\cdot)$ are extended to crumble contexts by:

$$\begin{aligned} \text{fv}(\langle \cdot \rangle) &:= \emptyset & \text{fv}((b, e[x \leftarrow \langle \cdot \rangle])) &:= \text{fv}((b, e)) \setminus \{x\} \\ \text{dom}(\langle \cdot \rangle) &:= \emptyset & \text{dom}((b, e[x \leftarrow \langle \cdot \rangle])) &:= \text{dom}(e) \cup \{x\} \end{aligned}$$

An environment context E is *well-named* if $E\langle x \rangle$ is well-named (for x fresh) and a crumble context C is *well-named* if $C := \langle \cdot \rangle$ or $C := (b, E)$ and E is well-named.

LEMMA B.4 (DECOMPOSITION OF READ BACK, AUXILIARY). Let e be an environment, and x be a variable such that $x \notin \text{dom}(e) \cup \text{fv}(e)$. Then $(c @ [x \leftarrow b]e)_{\downarrow} = (c @ e)_{\downarrow} \{x \leftarrow (b, e)_{\downarrow}\}$ for every crumble c and bite b .

PROOF. Let $c = (b', e')$. We proceed by structural induction on e :

- If $e = \epsilon$, then $(b, e)_{\downarrow} = b_{\downarrow}$ and hence $(b', e'[x \leftarrow b]e)_{\downarrow} = (b', e'[x \leftarrow b])_{\downarrow} = (b', e')_{\downarrow} \{x \leftarrow b_{\downarrow}\} = (b', e')_{\downarrow} \{x \leftarrow (b, e)_{\downarrow}\}$.
- Otherwise $e = e''[z \leftarrow b'']$ and then $(b, e)_{\downarrow} = (b, e'')_{\downarrow} \{z \leftarrow b'_{\downarrow}\}$, so

$$\begin{aligned} (b', e'[x \leftarrow b]e)_{\downarrow} &= (b', e'[x \leftarrow b']e''[z \leftarrow b''])_{\downarrow} \\ &= (b', e'[x \leftarrow b']e'')_{\downarrow} \{z \leftarrow b''_{\downarrow}\} \\ &= (b', e'e'')_{\downarrow} \{x \leftarrow (b', e'')_{\downarrow}\} \{z \leftarrow b''_{\downarrow}\} && \text{by i.h.} \\ &= (b', e'e'')_{\downarrow} \{z \leftarrow b''_{\downarrow}\} \{x \leftarrow (b', e'')_{\downarrow} \{z \leftarrow b''_{\downarrow}\}\} && \text{as } x \notin \text{fv}(b''_{\downarrow}) \\ &= (b', e'e)_{\downarrow} \{x \leftarrow (b', e)_{\downarrow}\} \end{aligned}$$

where $x \notin \text{fv}(b'_{\downarrow}) = \text{fv}(b'')$ by Remark 4.1 and the hypothesis that $x \notin \text{fv}(e)$. \square

LEMMA B.5 (READ-BACK VS. DISJOINTEDNESS). For every crumble c and environment e : if $c \# e$, then $(c @ e)_{\downarrow} = c_{\downarrow}$.

PROOF. By structural induction on e :

- If $e := \epsilon$ then $c @ e = c$ and hence $(c @ e)_{\downarrow} = c_{\downarrow}$.
- Otherwise $e := e'[x \leftarrow b]$. By *i.h.* (which can be applied since $c \# e$ implies $c \# e'$, because $\text{dom}(e') \subseteq \text{dom}(e)$), $(c @ e')_{\downarrow} = c_{\downarrow}$. From $c \# e$ it follows that $x \notin \text{fv}(c) = \text{fv}(c_{\downarrow})$ (by Remark 4.1), so $c_{\downarrow} \{x \leftarrow b_{\downarrow}\} = c_{\downarrow}$. Therefore, $(c @ e)_{\downarrow} = c @ e'_{\downarrow} \{x \leftarrow b_{\downarrow}\} = c_{\downarrow} \{x \leftarrow b_{\downarrow}\} = c_{\downarrow}$. \square

LEMMA B.6 (DECOMPOSITION OF READ BACK). Let c be a crumble and e be an environment such that $c \# e$. Then $(c @ [x \leftarrow b]e)_{\downarrow} = c_{\downarrow} \{x \leftarrow (b, e)_{\downarrow}\}$ for every bite b .

Note that Lemma B.6 does not hold without the hypothesis $c \# e$. Indeed, take $c := (y, \epsilon)$ and $e := [y \leftarrow zz]$ with $x \neq y$: for any term b , one has $(c @ [x \leftarrow b]e)_{\downarrow} = zz \neq y = c_{\downarrow}\{x \leftarrow (b, e)_{\downarrow}\}$.

PROOF. According to Lemma B.4,

$$(c @ [x \leftarrow b]e)_{\downarrow} = (c @ e)_{\downarrow}\{x \leftarrow (b, e)_{\downarrow}\} = c_{\downarrow}\{x \leftarrow (b, e)_{\downarrow}\}$$

where the last equality holds by Lemma B.5, since $c \# e$. \square

PROOF OF (PROP. 4.2). For every term t and every value v , one has $t_{\downarrow} = t$ and $\bar{v}_{\downarrow} = v$.

PROOF. By mutual induction on the term t and the value v . Cases:

- *Variable*, i.e. $t := x := v$; then, $\underline{t} = (x, \epsilon)$ and $\bar{v} = x$, thus $\underline{t}_{\downarrow} = x = t$ and $\bar{v}_{\downarrow} = x = v$.
- *Error or Boolean*: similar to the case *Variable* above.
- *Abstraction*, i.e. $t := \lambda x. u := v$; then, $\underline{t} = (\lambda x. \underline{u}, \epsilon)$ and $\bar{v} = \lambda x. \underline{u}$; by i.h., $\underline{u}_{\downarrow} = u$, hence $\underline{t}_{\downarrow} = \lambda x. \underline{u}_{\downarrow} = t$ and $\bar{v}_{\downarrow} = \lambda x. \underline{u}_{\downarrow} = v$.
- *Conditional*, case $t := (\text{if } v \text{ then } u \text{ else } s)$. Then $\underline{t} = (\text{if } \bar{v} \text{ then } \underline{u} \text{ else } \underline{s}, \epsilon)$. Since $\underline{t}_{\downarrow} = (\text{if } \bar{v}_{\downarrow} \text{ then } \underline{u}_{\downarrow} \text{ else } \underline{s}_{\downarrow})$, we can conclude by using the i.h.
- *Conditional*, case $t := (\text{if } u \text{ then } s \text{ else } r)$ with u not a value. Then $\underline{t} = (\text{if } x \text{ then } \underline{s} \text{ else } \underline{r}, \epsilon) @ [x \leftarrow q]e$ where $\underline{u} = (q, e)$. By i.h., $\underline{s}_{\downarrow} = s$ and $\underline{r}_{\downarrow} = r$. We can suppose that $(\{x\} \cup \text{fv}(s) \cup \text{fv}(r)) \cap \text{dom}(e) = \emptyset$ by the freshness condition in the definition of the transformation, and hence we can apply Lemma B.6 so that $\underline{t}_{\downarrow} = (\text{if } x \text{ then } \underline{s}_{\downarrow} \text{ else } \underline{r}_{\downarrow})\{x \leftarrow \underline{u}_{\downarrow}\} = t$.
- *Application of two values*, i.e. $t := vv'$; then, $\underline{t} = (\bar{v}\bar{v}', \epsilon)$; by i.h., $\bar{v}_{\downarrow} = v$ and $\bar{v}'_{\downarrow} = v'$, so $\underline{t}_{\downarrow} = \bar{v}_{\downarrow}\bar{v}'_{\downarrow} = t$.
- *Application of a non-value to a value*, i.e. $t := uv$ where u is not a value; then, $\underline{t} = (x\bar{v}, [x \leftarrow s]e) = (x\bar{v}, \epsilon) @ [x \leftarrow s]e$ where $\underline{u} = (s, e)$; we can suppose without loss of generality that $(\text{fv}(u) \cup \{x\}) \cap \text{dom}(e) = \emptyset$. By i.h., $\underline{u}_{\downarrow} = u$ and $\bar{v}_{\downarrow} = v$. By the freshness condition for x , we can apply Lemmas B.4 and B.5 so that $\underline{t}_{\downarrow} = (x\bar{v} @ e)_{\downarrow}\{x \leftarrow (s, e)_{\downarrow}\} = x\bar{v}_{\downarrow}\{x \leftarrow \underline{u}_{\downarrow}\} = t$.
- *Application of a term to a non-value*, i.e. $t := us$ where s is not a value; then, $\underline{t} = \underline{ux} @ ([x \leftarrow r]e)$ where $\underline{s} = (r, e)$ and $\underline{ux} = (yx, [y \leftarrow q]e') = (yx, \epsilon) @ ([y \leftarrow q]e')$ with $\underline{u} = (q, e')$; we can suppose without loss of generality that $(\text{fv}(u) \cup \{x\}) \cap \text{dom}(e) = \emptyset$ and $\{x, y\} \cap \text{dom}(e') = \emptyset$. By i.h., $\underline{u}_{\downarrow} = u$ and $\underline{s}_{\downarrow} = s$. By the freshness condition for x and y , we can apply Lemmas B.4 and B.5 so that $\underline{ux}_{\downarrow} = (\underline{yx} @ e')_{\downarrow}\{y \leftarrow (q, e')_{\downarrow}\} = yx\{y \leftarrow \underline{u}_{\downarrow}\} = ux$ and $\underline{t}_{\downarrow} = (\underline{ux} @ e)_{\downarrow}\{x \leftarrow (r, e)_{\downarrow}\} = \underline{ux}_{\downarrow}\{x \leftarrow \underline{s}_{\downarrow}\} = ux\{x \leftarrow s\} = t$. \square

LEMMA B.7. Let c, d be crumbles, and e be an environment. If $c_{\downarrow} = d_{\downarrow}$, then $(c @ e)_{\downarrow} = (d @ e)_{\downarrow}$.

PROOF. By induction on e :

- if $e := \epsilon$, then we conclude because $c @ e = c$ and $d @ e = d$;
- if $e := e'[x \leftarrow b]$, by i.h. $(c @ e')_{\downarrow} = (d @ e')_{\downarrow}$, and we conclude because $(c @ e)_{\downarrow} = (c @ e')_{\downarrow}\{x \leftarrow b_{\downarrow}\}$ and $(d @ e)_{\downarrow} = (d @ e')_{\downarrow}\{x \leftarrow b_{\downarrow}\}$. \square

PROPOSITION B.8. Let e be a crumbled environment. Then:

- (1) $(x v, e)_{\downarrow} = (x, e)_{\downarrow}(v, e)_{\downarrow}$ for every variable x and crumbled value v .

- (2) $(\text{if } x \text{ then } c \text{ else } d, e)_{\downarrow} = \text{if } (x, e)_{\downarrow} \text{ then } (c @ e)_{\downarrow} \text{ else } (d @ e)_{\downarrow}$ for every variable x and crumbles c, d .

PROOF. By induction on e :

- Case $e = \epsilon$:
 - (1) $(x v, \epsilon)_{\downarrow} = (x v)_{\downarrow} = x_{\downarrow} v_{\downarrow} = (x, \epsilon)_{\downarrow}(v, \epsilon)_{\downarrow}$.
 - (2) $(\text{if } x \text{ then } c \text{ else } d, \epsilon)_{\downarrow} = (\text{if } x \text{ then } c \text{ else } d)_{\downarrow} = (\text{if } x_{\downarrow} \text{ then } c_{\downarrow} \text{ else } d_{\downarrow})_{\downarrow} = (\text{if } ((x, \epsilon)_{\downarrow}) \text{ then } ((c @ \epsilon)_{\downarrow}) \text{ else } ((d @ \epsilon)_{\downarrow}))_{\downarrow}$.
- Case $e = e'[y \leftarrow b]$ for some e', y, b :
 - (1) By the definition of $(\cdot)_{\downarrow}$, $(x v, e'[y \leftarrow b])_{\downarrow} = (x v, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\}$. By i.h. $(x v, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\} = ((x, e')_{\downarrow}(v, e')_{\downarrow})\{y \leftarrow b_{\downarrow}\}$. By the definition of substitution, $((x, e')_{\downarrow}(v, e')_{\downarrow})\{y \leftarrow b_{\downarrow}\} = ((x, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\})(v, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\}$. Again by the definition of $(\cdot)_{\downarrow}$, $((x, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\})(v, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\} = (x, e)_{\downarrow}(v, e)_{\downarrow}$ and we conclude.
 - (2) By the definition of $(\cdot)_{\downarrow}$, $(\text{if } x \text{ then } c \text{ else } d, e'[y \leftarrow b])_{\downarrow} = (\text{if } x \text{ then } c \text{ else } d, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\}$. By i.h. $(\text{if } x \text{ then } c \text{ else } d, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\} = (\text{if } (x, e')_{\downarrow} \text{ then } (c @ e')_{\downarrow} \text{ else } (d @ e')_{\downarrow})\{y \leftarrow b_{\downarrow}\}$. By the definition of substitution, $(\text{if } (x, e')_{\downarrow} \text{ then } (c @ e')_{\downarrow} \text{ else } (d @ e')_{\downarrow})\{y \leftarrow b_{\downarrow}\} = (\text{if } (x, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\} \text{ then } (c @ e')_{\downarrow}\{y \leftarrow b_{\downarrow}\} \text{ else } (d @ e')_{\downarrow}\{y \leftarrow b_{\downarrow}\})_{\downarrow}$. Again by the definition of $(\cdot)_{\downarrow}$, $(\text{if } (x, e')_{\downarrow}\{y \leftarrow b_{\downarrow}\} \text{ then } (c @ e')_{\downarrow}\{y \leftarrow b_{\downarrow}\} \text{ else } (d @ e')_{\downarrow}\{y \leftarrow b_{\downarrow}\})_{\downarrow} = \text{if } (x, e)_{\downarrow} \text{ then } (c @ e)_{\downarrow} \text{ else } (d @ e)_{\downarrow}$ and we conclude. \square

LEMMA B.9. Let e be an environment such that $x \notin \text{dom}(e)$. Then $(x, e[x \leftarrow b]e')_{\downarrow} = (b, e')_{\downarrow}$ for every b, e' .

PROOF. By Lemma B.5, $(x, e)_{\downarrow} = (x, \epsilon)_{\downarrow}$ because $x \notin \text{dom}(e)$. By Lemma B.7, $(x, e[x \leftarrow b])_{\downarrow} = (x, [x \leftarrow b])_{\downarrow} = x\{x \leftarrow b_{\downarrow}\} = b_{\downarrow} = (b, \epsilon)_{\downarrow}$. Again by Lemma B.7, $(x, e[x \leftarrow b]e')_{\downarrow} = (b, e')_{\downarrow}$. \square

LEMMA B.10. For all crumble context C , crumble c , and environment e , one has $C\langle c \rangle @ e = C\langle c @ e \rangle$.

PROOF. By cases according to the definition of the crumble context C .

If $C := \langle \cdot \rangle$, then $C\langle c \rangle @ e = c @ e = C\langle c @ e \rangle$.

Otherwise $C := (b, e'[x \leftarrow \langle \cdot \rangle])$; let $c := (b', e'')$; then, $c @ e = (b', e''e)$ and hence $C\langle c \rangle @ e = (b, e'[x \leftarrow b']e'') @ e = (b, e'[x \leftarrow b']e''e) = (b, e'[x \leftarrow b']) @ e''e = C\langle c @ e \rangle$. \square

COROLLARY B.11 (READ BACK VS. CRUMBLE CONTEXTS). Let c be a crumble and C, C' be crumble contexts.

- (1) *Plugging*: If $C \# c$ and C_{\downarrow} is a λ -context and $C\langle c \rangle$ is well-named, then $C\langle c \rangle_{\downarrow} = C_{\downarrow}\langle c_{\downarrow} \rangle$.
- (2) *Composition*: If $C \# C'$ and $C\langle C' \rangle$ is well-named, where C_{\downarrow} and C'_{\downarrow} are contexts, then $C\langle C' \rangle_{\downarrow} = C_{\downarrow}\langle C'_{\downarrow} \rangle$, which is a context.

PROOF. (1) If $C := \langle \cdot \rangle$, then $C_{\downarrow} = \langle \cdot \rangle$ and so $C\langle c \rangle_{\downarrow} = c_{\downarrow} = C_{\downarrow}\langle c_{\downarrow} \rangle$. Otherwise $C := (b, e[x \leftarrow \langle \cdot \rangle])$ with $c = (b', e')$; since $C\langle c \rangle = (b, e[x \leftarrow b']e')$ is well-named, then $x \notin \text{dom}(e')$; therefore, from $C \# c$ it follows that $(b, e) \# e'$; we have $C\langle c \rangle_{\downarrow} = (b, e)_{\downarrow}\{x \leftarrow (b', e')_{\downarrow}\} = C_{\downarrow}\langle c_{\downarrow} \rangle$ by Lemma B.6 and because by hypothesis C unfolds to a context.

(2) The composition of two contexts is a context, thus $C_{\downarrow} \langle C'_{\downarrow} \rangle$ is a context since C_{\downarrow} and C'_{\downarrow} are so.

If $C := \langle \cdot \rangle$, then $C_{\downarrow} = \langle \cdot \rangle$ and $C \langle C' \rangle = C'$, thus $C \langle C' \rangle_{\downarrow} = C'_{\downarrow} = C_{\downarrow} \langle C'_{\downarrow} \rangle$.

If $C' := \langle \cdot \rangle$, then $C'_{\downarrow} = \langle \cdot \rangle$ and $C \langle C' \rangle = C$, so $C \langle C' \rangle_{\downarrow} = C_{\downarrow} = C_{\downarrow} \langle C'_{\downarrow} \rangle$.

Finally, if $C := (b, e[x \leftarrow \langle \cdot \rangle])$ and $C' = (b', e'[y \leftarrow \langle \cdot \rangle])$, then $C \langle C' \rangle = (b, e[x \leftarrow b']e'[y \leftarrow \langle \cdot \rangle])$ and so

$$\begin{aligned} C \langle C' \rangle_{\downarrow} &= (b, e[x \leftarrow b']e')_{\downarrow} \{y \leftarrow \langle \cdot \rangle\} \\ &= (b, e)_{\downarrow} \{x \leftarrow (b', e')_{\downarrow}\} \{y \leftarrow \langle \cdot \rangle\} && \text{by Lemma B.6} \\ &= (b, e)_{\downarrow} \{x \leftarrow (b', e')_{\downarrow}\} \{y \leftarrow \langle \cdot \rangle\} && \text{as } y \notin \text{fv}((b, e)_{\downarrow}) \\ &= C_{\downarrow} \langle C'_{\downarrow} \rangle \end{aligned}$$

where Lemma B.6 can be applied because $(b, e) \# e'$, since $C \# C'$ and $x \notin \text{dom}(e')$ (as $C \langle C' \rangle$ is well-named); and $y \notin \text{fv}((b, e)_{\downarrow})$ because $C \# C'$. \square

LEMMA B.12. *Let $c = C \langle (b, e) \rangle$ be a crumble. Then $\text{fv}(b), \text{fv}(C) \subseteq \text{dom}(e) \cup \text{fv}(c)$.*

PROOF. By cases according to the definition of the crumble context C .

If $C := \langle \cdot \rangle$ then $\text{fv}(C) = \emptyset \subseteq \text{dom}(e) \cup \text{fv}(c)$ and $c = (b, e)$, so $\text{fv}(c) = (\text{fv}(b) \setminus \text{dom}(e)) \cup \text{fv}(e)$ and hence $\text{fv}(b) \subseteq (\text{fv}(c) \setminus \text{fv}(e)) \cup \text{dom}(e) \subseteq \text{fv}(c) \cup \text{dom}(e)$.

Otherwise $C := (b, e'[x \leftarrow \langle \cdot \rangle])$ and then $\text{fv}(C) = \text{fv}(b') \cup (\text{dom}(e') \setminus \{x\})$ and $c = (u, e'[x \leftarrow b]e)$; therefore, $\text{fv}(c) = \text{fv}(C) \cup (\text{fv}(b) \setminus \text{dom}(e)) \cup \text{fv}(e)$ and hence $\text{fv}(C) \subseteq \text{dom}(e) \cup \text{fv}(c)$ and $\text{fv}(b) \subseteq \text{dom}(e) \cup \text{fv}(c)$. \square

Remark B.1. For every crumble c , one has $c @ [x \leftarrow \langle \cdot \rangle]_{\downarrow} = c_{\downarrow} \{x \leftarrow \langle \cdot \rangle\}$. Indeed, let $c := (b, e)$; then, $(b, e) @ [x \leftarrow \langle \cdot \rangle]_{\downarrow} = (b, e[x \leftarrow \langle \cdot \rangle])_{\downarrow} = (b, e)_{\downarrow} \{x \leftarrow \langle \cdot \rangle\}$.

PROOF OF (LEMMA 4.5). *For every term t :*

- (1) Freshness: \underline{t} is well-named.
- (2) Closure: if t is closed, then $\text{fv}(\underline{t}) = \emptyset$.
- (3) Disjointedness: $\text{dom}(C) \cap \text{fv}(b) = \emptyset$ if $\underline{t} = C \langle (b, e) \rangle$.
- (4) Bodies: every body in \underline{t} is the translation of a term.
- (5) Contextual decoding: if $\underline{t} = C \langle c \rangle$, then C_{\downarrow} is a right v-context.

PROOF.

- (1) It follows immediately from the freshness condition in the definition of translation.
- (2) By Remark 4.1.
- (3) It follows immediately from the freshness condition in the definition of translation.
- (4) By induction on t and by cases on the rules defining the translation.
- (5) By induction on the size of t . Cases:

- λ -Value, i.e. $t = v$. Then $\underline{t} = (\bar{v}, \epsilon)$ and so the only possible crumble context C such that $\underline{t} = C \langle c \rangle$ for some crumble c is $C = \langle \cdot \rangle$ and so $C_{\downarrow} = \langle \cdot \rangle$, which is a right v-context.
- λ -Value applied to value, i.e. $t = vv'$. As in the previous case.
- Application applied to a value, i.e. $t = uv$ with u not a λ -value. Then, $\underline{t} = (x\bar{v}, [x \leftarrow s]e)$ with $\underline{u} = (s, e)$ and x fresh. Cases for C :

– Empty, i.e. $C = \langle \cdot \rangle$: as for values (see above).

– Non-empty, i.e. $C = (x\bar{v}, [x \leftarrow \langle \cdot \rangle]) \langle C' \rangle$ where C' is a crumble context of \underline{u} , i.e. $\underline{u} = C' \langle c \rangle$. The read-back of the crumble context $C'' := (x\bar{v}, [x \leftarrow \langle \cdot \rangle])$ is $C''_{\downarrow} = x\bar{v}_{\downarrow} \{x \leftarrow \langle \cdot \rangle\} = \langle \cdot \rangle \bar{v}_{\downarrow}$, which is a right v-context because \bar{v}_{\downarrow} is a λ -value by Remark 4.1 val-to-val. By i.h., C'_{\downarrow} is a right v-context. By the freshness conditions in the definition of translation, $C'' \# C'$; according to Corollary B.11.2, $C_{\downarrow} = C'' \langle C' \rangle_{\downarrow} = C''_{\downarrow} \langle C'_{\downarrow} \rangle$, which is a right v-context since the composition of right v-context is a right v-context (Lemma A.1).

• λ -Term applied to application, i.e. $t = us$ with s not a λ -value. Then, $\underline{t} = \underline{ux} @ [x \leftarrow b]e$ where $\underline{s} = (b, e)$ and x is fresh. Cases for C :

– C is a crumble context of \underline{ux} , i.e. $\underline{ux} = C \langle c \rangle$, then C_{\downarrow} is a right v-context by i.h. (the size of the term ux is strictly less than the size of us because s is not a value).

– $C = C'' \langle C' \rangle$ where $C'' := \underline{ux} @ [x \leftarrow \langle \cdot \rangle]$ and C' is a crumble context of s i.e. $\underline{s} = C' \langle c \rangle$. By the freshness condition in the definition of translation and according to Remark B.1, $C''_{\downarrow} = \underline{ux}_{\downarrow} \{x \leftarrow \langle \cdot \rangle\} = \underline{u}_{\downarrow} \langle \cdot \rangle$, which is a right v-context. By i.h., C'_{\downarrow} is a right v-context. According to Corollary B.11.2, $C_{\downarrow} = C'' \langle C' \rangle_{\downarrow} = C''_{\downarrow} \langle C'_{\downarrow} \rangle$, which is a right v-context since the composition of right v-contexts is a right v-context (Lemma A.1).

• Conditional, case $t := (\text{if } v \text{ then } u \text{ else } s)$. Then $\underline{t} = (\text{if } \bar{v} \text{ then } \underline{u} \text{ else } \underline{s}, \epsilon)$. Necessarily $C = \langle \cdot \rangle$, and hence C_{\downarrow} is a right v-context.

• Conditional, case $t := (\text{if } u \text{ then } s \text{ else } r)$ with u not a value. Then $\underline{t} = (\text{if } x \text{ then } \underline{s} \text{ else } \underline{r}, [x \leftarrow q]e)$ where $\underline{u} = (q, e)$. We proceed by cases on C :

– Case $C = \langle \cdot \rangle$. Then clearly C_{\downarrow} is a right v-context.

– Case $C = C'' \langle C' \rangle$, where $C'' := (\text{if } x \text{ then } \underline{s} \text{ else } \underline{r}, [x \leftarrow \langle \cdot \rangle])$ and C' is a crumble context of \underline{u} . By the freshness condition in the definition of translation and according to Remark B.1, $C''_{\downarrow} = (\text{if } \langle \cdot \rangle \text{ then } \underline{s} \text{ else } \underline{r})$, which is a right v-context. By i.h., C'_{\downarrow} is a right v-context. According to Corollary B.11.2, $C_{\downarrow} = C'' \langle C' \rangle_{\downarrow} = C''_{\downarrow} \langle C'_{\downarrow} \rangle = (\text{if } C'_{\downarrow} \text{ then } \underline{s} \text{ else } \underline{r})$, which is a right v-context. \square

C PROOFS OF SECT. 5 (CLOSED CASE)

C.1 Proofs of Subsect. 5.1

LEMMA C.1 (CLOSURE UNDER SUBSTITUTION). *Let v and v' be values. Then $v\{x \leftarrow v'\}$ is a value. If moreover v is a λ -abstraction, then $v\{x \leftarrow v'\}$ is so.*

PROOF. By cases on the definition of value:

- Variable, i.e. either $v := x$ and then $v\{x \leftarrow v'\} = v'$ which is a value by hypothesis; or $v := y \neq x$ and then $v\{x \leftarrow v'\} = y$ which is a value.
- Abstraction, i.e. $v := \lambda y. u$ and we can suppose without loss of generality that $y \notin \text{fv}(v) \cup \{x\}$; therefore, $v\{x \leftarrow v'\} = \lambda y. (u\{x \leftarrow v'\})$ which is a λ -abstraction and hence a value.

- *Booleans and errors*, i.e. $v \in \{\text{true}, \text{false}, \text{err}\}$. Trivial since $v\{x \leftarrow v'\} = v$. \square

LEMMA C.2 (READ-BACK TO VALUE). *For every crumbled value v and v -environment e_v , one has that $(v, e_v)_\downarrow$ is a value. If moreover v is an abstraction, then $(v, e_v)_\downarrow$ is a λ -abstraction.*

PROOF. By induction on the length of e_v . We proceed by cases on the shape of v .

- *Abstraction*: If $v := \epsilon$, then clearly $(v, e_v)_\downarrow = v_\downarrow$ is a λ -abstraction and hence a value. Otherwise $v := e'_v[x \leftarrow b]$ where b is a crumbled value (and hence b_\downarrow is a value): thus, $(v, e_v)_\downarrow = (v, e'_v)_\downarrow\{x \leftarrow b_\downarrow\}$; by i.h., $(v, e'_v)_\downarrow$ is a λ -abstraction, thus $(v, e_v)_\downarrow$ is a λ -abstraction (and so a value) by Lemma C.1.
- *Booleans and errors*: the proof is identical to the previous case.
- *Variable*: If $v := x \notin \text{dom}(e_v)$, then $(v, e_v)_\downarrow = x$ which is a value; otherwise $v := x \in \text{dom}(e_v)$ with $e_v := e'_v[x \leftarrow \lambda y.c]e''_v$, and then $(x, e_v)_\downarrow = (\lambda y.c, e''_v)_\downarrow$ (since $x \notin \text{dom}(e''_v)$) is a value by i.h., according to the previous point. \square

LEMMA C.3. *Let $c = (b, e_v)$ be a well-named closed crumble, and b have the following property: b is a value, or b is x or xv or if x then c else d but x is not defined in e_v . Then c is a v -crumble.*

PROOF. Let $c = (b, e_v)$ as above: it suffices to prove that b is a value. This follows easily from the hypothesis that c is closed, since the cases where b is x or xv or if x then c else d but x is not defined in e_v are impossible. \square

COROLLARY C.4. *If the well-named closed crumble (b, e_v) is normal, then it is a v -crumble.*

PROOF OF (PROP. 5.3). A closed crumble c is normal if and only if it is a v -crumble.

PROOF.

- (\Rightarrow) Let $c = (b, e)$ be well-named, closed, and normal. We proceed by structural induction on e :
- if $e = \epsilon$, then (b, e) is a v -crumble by Cor. C.4;
 - if $e = [x \leftarrow b']e'$, then also the crumble (b', e') is normal. By i.h. (b', e') is a v -crumble, and therefore $e = [x \leftarrow b']e'$ is a v -environment. By Cor. C.4, $c = b, e$ is a v -crumble.
- (\Leftarrow) Let $c = c_v$, we need to prove that c_v is normal. Let $c_v = C\langle(v_{\neg x}, e_v)\rangle$ for some $C, v_{\neg x}, e_v$. Clearly no reduction rule is applicable, because $v_{\neg x}$ is not a variable or an application. \square

C.2 Proofs of Subsect. 5.2

PROOF OF (LEMMA 5.5). For every reachable crumble c in the Crumble GLAM:

- (1) Freshness: c is well-named.
- (2) Closure: $\text{fv}(c) = \emptyset$.
- (3) Bodies: every body occurring in c is a subterm (up to renaming) of the initial crumble.
- (4) Weak contextual decoding: for every decomposition $C\langle(b, e_v)\rangle$ where b is not a crumbled value, if C'' is a prefix of C then C''_\downarrow is a right v -context.

PROOF. By induction on the length of the reduction sequence leading to the crumble. The base cases hold by Lemma 4.5 (by noting that for Point 4, Lemma 4.5.5 implies the weaker statement Lemma 5.5.4). As for the inductive cases, we inspect each transition:

1. The substitution transitions $\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}$ do not change the set of variables occurring on the lhs of substitutions outside abstractions because they copy a value that does not contain any. Hence the claim follows from the i.h.. For transition β_v the claim follows from the side condition. For the remaining rules $\text{ift}, \text{iff}, \text{ife}, @e$ the claim follows from the fact that all substitutions outside abstractions in the rhs already occur in the lhs.
2. The substitution transitions $\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}$ do not change the domain of the crumble and only copy to the left a value from the environment, and the claim follows from the i.h.. Transition β_v copies to the top level and renames the body of an abstraction. By the properties of α -renaming $\text{fv}((c @ [x \leftarrow v])^\alpha) = \text{fv}(c @ [x \leftarrow v]) = \text{fv}(\lambda x.c)$, and since by i.h. $\text{fv}(\lambda x.c) \subseteq \text{dom}(e_v)$, we can conclude with $\text{fv}((c @ [x \leftarrow v])^\alpha) \subseteq \text{dom}(e_v)$. The terms in the rhs of the remaining transitions $\text{ift}, \text{iff}, \text{ife}, @e$ already occur in the lhs under the same environment. Therefore the claim follows from the i.h..
3. The rules $\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}$ may copy an abstraction, but the abstraction was already in the environment, and the claim follows from the i.h.. The rule β_v copies and renames the body of an abstraction that was already in the environment, and the claim follows from the i.h. since the translation commutes with the renaming of free variables (Remark 4.1.3). All the bodies in the rhs of the remaining rules $\text{ift}, \text{iff}, \text{ife}, @e$ already occur in the lhs and therefore the claim follows from the i.h..
4. Let $\underline{b''} \rightarrow^n C'\langle(b', e'_v)\rangle \rightarrow_a C\langle(b, e_v)\rangle$ (where b is not a practical value). Cases of the reduction step $C'\langle(b', e'_v)\rangle \rightarrow_a C\langle(b, e_v)\rangle$:
 - Case β_v : $C'\langle((\lambda x.c)v, e_v)\rangle \rightarrow_{\beta_v} C'\langle c^\alpha @ ([x \leftarrow v]e_v)\rangle$. Let C'' be a prefix of C . There are two sub-cases:
 - * C'' is a prefix of C' : by i.h. C''_\downarrow is a right v -context.
 - * C' is a prefix of C'' , i.e. $C'' = C'\langle C''' \rangle$ and $c^\alpha = C''' \langle c' \rangle$. By Lemma 4.5.4 and Lemma 5.5.3 c is the translation of a λ -term, by Remark 4.1.3 c^α is so, and thus by Lemma 4.5.5 C''_\downarrow is a right v -context. By i.h., C'_\downarrow is a right v -context as well. Since $C''_\downarrow = C'_\downarrow \langle C''' \rangle$ according to Corollary B.11.2, we obtain that C''_\downarrow is a right v -context as composition of right v -contexts (Lemma A.1).
 - Case ift : $C'\langle(\text{if true then } c \text{ else } d, e_v)\rangle \rightarrow_{\beta_v} C'\langle c @ e_v \rangle$. Let C'' be a prefix of C . There are two sub-cases:
 - * C'' is a prefix of C' : by i.h. C''_\downarrow is a right v -context.
 - * C' is a prefix of C'' , i.e. $C'' = C'\langle C''' \rangle$ and $c = C''' \langle c' \rangle$. By Lemma 4.5.4 and Lemma 5.5.3 c is the translation of a λ -term and thus by Lemma 4.5.5 C''_\downarrow is a right v -context. By i.h., C'_\downarrow is a right v -context as well. Since $C''_\downarrow = C'_\downarrow \langle C''' \rangle$ according to Corollary B.11.2, we obtain that C''_\downarrow is a right v -context as composition of right v -contexts (Lemma A.1).

- Case iff: identical to the previous case.
- Cases sub_{var} , ife, @e: they follow from the i.h. since C is necessarily a prefix of C' because b is a practical value.
- Cases sub_l and sub_{if} : they follow from the i.h., since $e'_v = e_v$ and $C = C'$. \square

LEMMA C.5 (DETERMINISM). \rightarrow_{C_r} is deterministic.

PROOF. Assume that there exists a crumble that may be decomposed in two ways $C\langle(b, e_v)\rangle = C'\langle(b', e'_v)\rangle$ such that they reduce respectively $C\langle(b, e_v)\rangle \rightarrow_a C\langle c\rangle$ and $C'\langle(b', e'_v)\rangle \rightarrow_b C'\langle d\rangle$ with rules $a, b \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e, \text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$.

We prove that it must necessarily be $a = b$, $C = C'$, and $c = d$ (up to alpha). Three cases:

- C strict initial segment of C' , i.e. $C' = C\langle C''\rangle$ for some $C'' \neq \langle \cdot \rangle$. We show that this case is not possible: in fact, it follows that $e_v = E\langle(b', e'_v)\rangle$ for some E , thus (b', e'_v) is a v -crumble, and by Prop. 5.3 it must be normal, contradicting the hypothesis that (b', e'_v) and c reduce with rule b .
- $C = C'$. By inspection of the reduction rules, $a = b$: in fact the rule β_v applies only when b is the application of an abstraction to a crumbled value, the rule sub_{var} only when b is a variable, and the rule sub_l only when b is the application of a variable to a crumbled value, etc. It remains to show that $c = d$ (up to alpha): this follows from the determinism of the lookup in the environment during sub_{var} , sub_l and sub_{if} reductions.
- C' initial segment of C , i.e. $C = C'\langle C''\rangle$. Symmetric to the first case. \square

LEMMA C.6. In every reachable crumble $C\langle(b, e)\rangle$ one has $\text{fv}(b) \cap \text{dom}(C) = \emptyset$.

PROOF. By Lemma B.12 and Lemmas 5.5.1-2. \square

PROPOSITION C.7 (OVERHEAD TRANSPARENCY). Let c be a reachable crumble, and let $a \in \{\text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$. If $c \rightarrow_a d$ then $c_\downarrow \rightarrow_a d_\downarrow$.

PROOF. Let $c := C\langle(b, e_v)\rangle \rightarrow_a C\langle(b', e'_v)\rangle =: d$, and let e'_v, e''_v such that $e_v = e'_v[x \leftarrow e_v(x)]e''_v$, noting that x does not occur in e''_v by Lemma 5.5.1 and Lemma C.6. We first prove that $(b, e_v)_\downarrow = (b', e'_v)_\downarrow$:

- Case sub_{var} , i.e. $b := x$ and $b' = e_v(x)$. By Lemma B.4, $(x, e'_v[x \leftarrow e_v(x)]e''_v)_\downarrow = (x, e'_v e''_v)_\downarrow \{x \leftarrow (e_v(x), e''_v)_\downarrow\} = (e_v(x), e''_v)_\downarrow$ as c is well-named (Lemma 5.5.1). By Lemma C.6, $\text{fv}(e_v(x)) \cap \text{dom}(e'_v[x \leftarrow e_v(x)]) = \emptyset$, therefore $(e_v(x), e''_v)_\downarrow = (e_v(x), e_v)_\downarrow$, and we conclude with $(x, e_v)_\downarrow = (e_v(x), e_v)_\downarrow$.
- Case sub_l , i.e. $b := x v$ and $b' = e_v(x) v$. Since $(x v, e_v)_\downarrow = (x, e_v)_\downarrow (v, e_v)_\downarrow$ (Prop. B.8), we can use the point above to conclude.
- Case sub_{if} , i.e. $b := \text{if } x \text{ then } c \text{ else } d$ and $b' = \text{if } e_v(x) \text{ then } c \text{ else } d$. Since $(\text{if } x \text{ then } c \text{ else } d, e_v)_\downarrow = \text{if } (x, e_v)_\downarrow \text{ then } (c @ e_v)_\downarrow \text{ else } (d @ e_v)_\downarrow$ (Prop. B.8), we can use the point above to conclude.

We now prove that $C\langle(b, e_v)\rangle_\downarrow = C\langle(b', e'_v)\rangle_\downarrow$ under the hypothesis that $(b, e_v)_\downarrow = (b', e'_v)_\downarrow$. By cases on C : if $C := \langle \cdot \rangle$ just use the hypothesis. Otherwise $C :=$

$(b'', e[x \leftarrow \langle \cdot \rangle])$ and so $(b'', e[x \leftarrow b]e_v)_\downarrow = (b'', ee_v)_\downarrow \{x \leftarrow (b, e_v)_\downarrow\} = (b'', ee_v)_\downarrow \{x \leftarrow (b', e'_v)_\downarrow\} = (b'', e[x \leftarrow b']e_v)_\downarrow$ by Lemma B.4. \square

LEMMA C.8 (SUBSTITUTION). Let t and u be terms, and v be a value. If $t \rightarrow_{\text{pif}} u$ then $t\{x \leftarrow v\} \rightarrow_{\text{pif}} u\{x \leftarrow v\}$.

PROOF. By induction on the definition of $t \rightarrow_{\text{pif}} u$. Cases:

- Root-step.
 - β_v -step, i.e. $t := (\lambda y.s)v' \mapsto_{\beta_v} s\{y \leftarrow v'\} =: u$ and we can suppose without loss of generality that $y \notin \text{fv}(v) \cup \{x\}$. According to Lemma C.1, $v'\{x \leftarrow v\}$ is a value. As a consequence, $t\{x \leftarrow v\} = (\lambda y.s\{x \leftarrow v\})(v'\{x \leftarrow v\}) \rightarrow_{\beta_v} s\{y \leftarrow v'\{x \leftarrow v\}\} = s\{y \leftarrow v'\}\{x \leftarrow v\} = u\{x \leftarrow v\}$.
 - the ift, iff, ife, @e steps are similar to the β_v -step. Lemma C.1 is used in the proof of the ife-step to prove that a ife-redex where the guard is an abstraction is mapped to a ife-redex of the same kind.
- Application right, i.e. $t := sr \rightarrow_{\text{pif}} sq =: u$ with $r \rightarrow_{\text{pif}} q$; by i.h. $r\{x \leftarrow v\} \rightarrow_{\text{pif}} q\{x \leftarrow v\}$, and therefore $t\{x \leftarrow v\} = s\{x \leftarrow v\}(r\{x \leftarrow v\}) \rightarrow_{\text{pif}} s\{x \leftarrow v\}(q\{x \leftarrow v\}) = u\{x \leftarrow v\}$.
- Application left, i.e. $t := sv' \rightarrow_{\text{pif}} rv' =: u$ with $s \rightarrow_{\text{pif}} r$; by i.h., $s\{x \leftarrow v\} \rightarrow_{\text{pif}} r\{x \leftarrow v\}$; according to Lemma C.1, $v'\{x \leftarrow v\}$ is a value and hence $t\{x \leftarrow v\} = s\{x \leftarrow v\}(v'\{x \leftarrow v\}) \rightarrow_{\text{pif}} r\{x \leftarrow v\}(v'\{x \leftarrow v\}) = u\{x \leftarrow v\}$.
- If-then-else guard, i.e. $t := \text{if } r \text{ then } s \text{ else } p \rightarrow_{\text{pif}} \text{if } q \text{ then } s \text{ else } p =: u$ with $r \rightarrow_{\text{pif}} q$; by i.h. $r\{x \leftarrow v\} \rightarrow_{\text{pif}} q\{x \leftarrow v\}$, and therefore $t\{x \leftarrow v\} = \text{if } r\{x \leftarrow v\} \text{ then } s\{x \leftarrow v\} \text{ else } p\{x \leftarrow v\} \rightarrow_{\text{pif}} \text{if } q\{x \leftarrow v\} \text{ then } s\{x \leftarrow v\} \text{ else } p\{x \leftarrow v\} = u\{x \leftarrow v\}$. \square

LEMMA C.9. Let c, d be crumbles, and let e_v be a v -environment. If $c_\downarrow \rightarrow_{\text{pif}} d_\downarrow$, then $(c @ e_v)_\downarrow \rightarrow_{\text{pif}} (d @ e_v)_\downarrow$.

PROOF. By induction on the length of e_v . If $e_v := \epsilon$ then $(c @ e_v)_\downarrow = c_\downarrow \rightarrow_{\text{pif}} d_\downarrow = (d @ e_v)_\downarrow$. Otherwise $e_v := e'_v[x \leftarrow v]$ where v is a practical value (and hence v_\downarrow is a value); by i.h., $(c @ e'_v)_\downarrow \rightarrow_{\text{pif}} (d @ e'_v)_\downarrow$ and hence $(c @ e_v)_\downarrow = (c @ e'_v)_\downarrow \{x \leftarrow v_\downarrow\} \rightarrow_{\text{pif}} (d @ e'_v)_\downarrow \{x \leftarrow v_\downarrow\} = (d @ e_v)_\downarrow$ according to Lemma C.8. \square

PROPOSITION C.10 (PRINCIPAL PROJECTION). Let c be a reachable crumble. If $c \rightarrow_a d$ for $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e\}$ then $c_\downarrow \rightarrow_a d_\downarrow$.

PROOF. Note that for every b, e_v , $(b, e_v) = (b, \epsilon) @ e_v$. Therefore all steps can be written in the form $C\langle(b, e_v)\rangle \rightarrow_a C\langle(c @ e_v)\rangle$ where b is not a crumbled value. The crumble context C unfolds to a right v -context by Lemma 5.5.4. We need to prove that $C\langle(b, e_v)\rangle_\downarrow \rightarrow_a C\langle(c @ e_v)\rangle_\downarrow$. By Lemma B.10 and Lemma C.9, it suffices to prove that $C\langle(b, \epsilon)\rangle_\downarrow \rightarrow_a C\langle c\rangle_\downarrow$.

We proceed by cases on the rule \rightarrow_a .

- Rule β_v : we need to prove that $C\langle((\lambda x.c)v, \epsilon)\rangle_{\downarrow} \rightarrow_{\beta_v} C\langle e @ [y \leftarrow v]\rangle_{\downarrow}$ where $e @ [y \leftarrow v] := (c @ [x \leftarrow v])^\alpha$.

$$\begin{aligned}
C\langle((\lambda x.c)v, \epsilon)\rangle_{\downarrow} &= C_{\downarrow}\langle(\lambda x.c_{\downarrow})v_{\downarrow}\rangle && \text{by Corollary B.11.1} \\
&=_{\alpha} C_{\downarrow}\langle(\lambda y.e_{\downarrow})v_{\downarrow}\rangle \\
&\rightarrow_{\beta_v} C_{\downarrow}\langle e_{\downarrow}\{y \leftarrow v_{\downarrow}\}\rangle \\
&= C_{\downarrow}\langle e @ [y \leftarrow v]\rangle_{\downarrow} \\
&= C\langle e @ [y \leftarrow v]\rangle_{\downarrow} && \text{by Corollary B.11.1} \\
&= C\langle(c @ [x \leftarrow v])^\alpha\rangle_{\downarrow}.
\end{aligned}$$

Note that the second use of Corollary B.11.1 requires that $C \# (e @ [y \leftarrow v])$ i.e. that $\text{fv}(C) \cap \text{dom}(e @ [y \leftarrow v]) = \emptyset$, which follows from the side condition about α -renaming in the β_v rule.

- Rules ift, iff, ife, @e: a quick check by cases over $(b, e_v) \mapsto_a c @ e_v$ shows that $b_{\downarrow} \rightarrow_a c_{\downarrow}$.

For example, $(\text{if true then } c \text{ else } d, e_v) \mapsto_{\text{ift}} c @ e_v$ and $(\text{if true then } c \text{ else } d)_{\downarrow} = \text{if true then } c_{\downarrow} \text{ else } d_{\downarrow} \mapsto_{\text{ift}} c_{\downarrow}$. The other cases are all similar.

Thus

$$\begin{aligned}
C\langle(b, \epsilon)\rangle_{\downarrow} &= C_{\downarrow}\langle b_{\downarrow}\rangle && \text{by Corollary B.11.1} \\
&\rightarrow_a C_{\downarrow}\langle c_{\downarrow}\rangle \\
&= C\langle c\rangle_{\downarrow} && \text{by Corollary B.11.1}
\end{aligned}$$

Note that the second use of Corollary B.11.1 requires that $C \# c$. The property holds because all substitutions outside abstractions in the rhs of the rules \rightarrow_a under consideration were such in the lhs, and because reachable crumbs are well-named (Lemma 5.5.1). \square

LEMMA C.11 (HALT). *Let c be a closed crumble. If c is Cr-normal then c_{\downarrow} is Pif-normal.*

PROOF. By Prop. 5.3, if c is normal then it is a v -crumble i.e. $c = c_v$. By Lemma C.2, c_{\downarrow} is a value. By Pif harmony (Prop. 2.1), c_{\downarrow} is \rightarrow_{pif} -normal. \square

PROOF OF (THM. 5.6). *Let c be a crumble reachable by the Crumble GLAM.*

- (1) Initialization: $t_{\downarrow} = t$ for every term t .
- (2) Principal projection: if $c \rightarrow_a d$ then $c_{\downarrow} \rightarrow_a d_{\downarrow}$, for any rule $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e\}$.
- (3) Overhead transparency: if $c \rightarrow_a d$ then $c_{\downarrow} = d_{\downarrow}$ for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$.
- (4) Determinism: the transition \rightarrow_{C_r} is deterministic.
- (5) Halt: if c is \rightarrow_{C_r} -normal then c_{\downarrow} is \rightarrow_{pif} -normal.
- (6) Overhead termination: \rightarrow_a terminates, for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$.

Therefore, the Crumble GLAM, Pif evaluation \rightarrow_{pif} , and the read-back $(\cdot)_{\downarrow}$ form an implementation system.

PROOF. (1) See Prop. 4.2.

(2) See Prop. C.10.

(3) See Prop. C.7.

(4) See Lemma C.5.

(5) See Lemma C.11, since c is closed by Lemma 5.5.2.

- (6) Immediate consequence of Lemma 5.8 (proved independently). \square

C.3 Proofs of Subsect. 5.3

PROOF OF (PROP. 5.7). *Let t be a term and v a value. Then:*

- (1) $|t|_{\text{var}} \leq 1$; and $|t|_{\text{var}} = 1$ if and only if t is a variable;
- (2) $|\bar{v}|_{\text{var}} \leq 1$; and $|\bar{v}|_{\text{var}} = 1$ if and only if v is a variable.

PROOF. By mutual induction on t and v :

- *Variable*, i.e. $v = t := x$: then $\bar{v} = x$ and $\underline{t} = (x, \epsilon)$, hence $|\bar{v}|_{\text{var}} = |\underline{t}|_{\text{var}} = 1$.
- *Boolean or error*, i.e. $v = t \in \{\text{true}, \text{false}, \text{err}\}$: then $\bar{v} = v$ and $\underline{t} = (v, \epsilon)$, hence $|\bar{v}|_{\text{var}} = |\underline{t}|_{\text{var}} = 0$.
- *Abstraction*, i.e. $v := \lambda x.u := t$: then $\bar{v} = \lambda x.\bar{u}$ and $\underline{t} = (\lambda x.\underline{u}, \epsilon)$; by i.h., $|\underline{u}|_{\text{var}} \leq 1$. We have $|\bar{v}|_{\text{var}} = |\underline{t}|_{\text{var}} = 0$ by definition.
- *Application of two λ -values*, i.e. $t := v'v''$: then, $\underline{t} = (\bar{v}'\bar{v}'', \epsilon)$. We have $|\underline{t}|_{\text{var}} = 0$ by definition.
- *Conditional on a value*, i.e. $t := \text{if } v' \text{ then } u \text{ else } s$: then, $\underline{t} = (\text{if } \bar{v}' \text{ then } \underline{u} \text{ else } \underline{s}, \epsilon)$. We have $|\underline{t}|_{\text{var}} = 0$ by definition.
- *Application of a non- λ -value to a λ -value*, i.e. $t := uv$ where u is not a value: then, $\underline{t} = (x\bar{v}, [x \leftarrow b]e)$ where $\underline{u} = (b, e)$; by i.h., $|\underline{u}|_{\text{var}} = 0$. Note that $|\underline{t}|_{\text{var}} = |(x\bar{v}, [x \leftarrow b]e)|_{\text{var}} = |b|_{\text{var}} + |e|_{\text{var}} = |(b, e)|_{\text{var}} = |\underline{u}|_{\text{var}} =_{\text{i.h.}} 0$.
- *Application of a term to a non-value*, i.e. $t := us$ where s is not a value: then, $\underline{t} = \underline{ux} @ ([x \leftarrow b]e) = (yx, [y \leftarrow b']e'[x \leftarrow b]e)$ where $\underline{s} = (b, e)$ and $\underline{ux} = (yx, [y \leftarrow b']e')$ with $\underline{u} = (b', e')$. By i.h., $|\underline{ux}|_{\text{var}} = |\underline{s}|_{\text{var}} = 0$. Since $|\underline{t}|_{\text{var}} = |\underline{u}|_{\text{var}} + |\underline{s}|_{\text{var}}$, we have $|\underline{t}|_{\text{var}} = 0$.
- *Conditional on a non-value*, i.e. $t := \text{if } u \text{ then } s \text{ else } r$ where u is not a value: then, $\underline{t} = (\text{if } x \text{ then } \underline{s} \text{ else } \underline{r}, [x \leftarrow b]e)$ with $\underline{u} = (b, e)$. By i.h., $|\underline{u}|_{\text{var}} = 0$. Then

$$|\underline{t}|_{\text{var}} = |(\text{if } x \text{ then } \underline{s} \text{ else } \underline{r}, [x \leftarrow b]e)|_{\text{var}} = |b|_{\text{var}} + |e|_{\text{var}} = |(b, e)|_{\text{var}} = |\underline{u}|_{\text{var}} =_{\text{i.h.}} 0$$

\square

We define the size $|\cdot|$ of crumbled forms, crumbs and environments as expected:

$$\begin{aligned}
|(b, e)| &:= |b| + |e| & |\epsilon| &:= 0 & |e[x \leftarrow b]| &:= 1 + |e| + |b| \\
|x| &:= 1 & |\lambda x.c| &:= 1 + |c| & |vw| &:= 1 + |v| + |w| \\
|\text{true}| &:= 1 & |\text{false}| &:= 1 & |\text{err}| &:= 1 \\
|\text{if } v \text{ then } c \text{ else } d| &:= 1 + |v| + |c| + |d|
\end{aligned}$$

PROOF OF (LEMMA 5.9). *Let t be a term and v a value. Then $|t| \leq 5|t|$ and $|\bar{v}| \leq 5|v|$.*

PROOF. By mutual induction on t and v :

- *Variable*, i.e. $v := x := t$: then $\bar{v} = x$ and $\underline{t} = (x, \epsilon)$, hence $|\bar{v}| = 1 \leq 5|v|$ and $|\underline{t}| = |x| + |\epsilon| = 1 \leq 5|t|$.
- *Error or Boolean*: similar to the case above.
- *Abstraction*, i.e. $v := \lambda x.u := t$: then $\bar{v} = \lambda x.\bar{u}$ and $\underline{t} = (\lambda x.\underline{u}, \epsilon)$; by i.h., $|\underline{u}| \leq 5|u|$ and hence $|\bar{v}| = |\underline{u}| + 1 \leq 5|u| + 1 \leq 5(|u| + 1) = 5|v|$ and $|\underline{t}| = |\lambda x.\underline{u}| + |\epsilon| = |\underline{u}| + 1 \leq 5|u| + 1 \leq 5(|u| + 1) = 5|t|$.
- *Application of two λ -values*, i.e. $t := vv'$: then, $\underline{t} = (\bar{v}\bar{v}', \epsilon)$; by i.h., $|\bar{v}| \leq 5|v|$ and $|\bar{v}'| \leq 5|v'|$; hence $|\underline{t}| = |\bar{v}\bar{v}'| + |\epsilon| = |\bar{v}| + |\bar{v}'| + 1 \leq 5|v| + 5|v'| + 1 \leq 5(|v| + |v'| + 1) = 5|t|$.

- *Application of a non- λ -value to a λ -value*, i.e. $t := uv$ where u is not a value: then, $\underline{t} = (x\bar{v}, [x \leftarrow b]e)$ where $\underline{u} = (b, e)$; by i.h., $|\bar{v}| \leq 5|v|$ and $|\underline{u}| = |b| + |e| \leq 5|u|$; hence $|\underline{t}| = |x| + |\bar{v}| + 1 + |b| + |e| + 1 = |\bar{v}| + |\underline{u}| + 3 \leq 5|v| + 5|u| + 3 \leq 5(|v| + |u| + 1) = 5|t|$.
- *Application of a term to a non-value*, i.e. $t := us$ where s is not a value: then, $\underline{t} = \underline{ux} @ ([x \leftarrow b]e) = (yx, [y \leftarrow b']e'[x \leftarrow b]e)$ where $\underline{s} = (b, e)$ and $\underline{ux} = (yx, [y \leftarrow b']e')$ with $\underline{u} = (b', e')$. By i.h., $|\underline{s}| = |b| + |e| \leq 5|s|$ and $|\underline{ux}| = |b'| + |e'| \leq 5|u|$. Therefore, $|\underline{t}| = |yx| + |b'| + |e'| + 1 + |b| + |e| + 1 = |\underline{u}| + |\underline{s}| + 5 \leq 5|u| + 5|s| + 5 \leq 5(|u| + |s| + 1) = 5|t|$.
- *Conditional*, case $t := (\text{if } v \text{ then } u \text{ else } s)$. Then $\underline{t} = (\text{if } \bar{v} \text{ then } \underline{u} \text{ else } \underline{s}, \epsilon)$ and $|\underline{t}| = 1 + |\bar{v}| + |\underline{u}| + |\underline{s}|$. Conclude by i.h.
- *Conditional*, case $t := (\text{if } u \text{ then } s \text{ else } r)$ with u not a value. Then $\underline{t} = (\text{if } x \text{ then } \underline{s} \text{ else } \underline{r}, [x \leftarrow q]e)$ where $\underline{u} = (q, e)$. Then $|\underline{t}| = 1 + 1 + |\underline{s}| + |\underline{r}| + 1 + |\underline{u}|$. Conclude by i.h. \square

PROOF OF (THM. 5.10). For any closed term t and any Crumble GLAM execution $\rho: \underline{t} \rightarrow_{Cr}^* c$, the cost of implementing ρ on a RAM is $O((|\rho|_p + 1) \cdot |t|)$ plus the cost of plugging and unplugging.

PROOF. By Lemma 5.8 and the discussion in the body about costs. \square

C.4 The Pointed Crumble GLAM

As discussed at the beginning of Sect. 5, we call the Crumble GLAM a machine even though it does not satisfy the usual requirements for abstract machines. One of these requirements is the presence of rules that guide evaluation to the next redex: in the Crumble GLAM the search for the next redex is instead left implicit in the evaluation rules, because it corresponds to going through the environment from right to left.

In order to implement the Crumble GLAM while respecting the complexity analysis that we just presented in Subsect. 5.3, we introduce a variant of the Crumble GLAM called Pointed Crumble GLAM, where the search for the next redex is decomposed into $O(1)$ steps of a new *search* transition. The other steps of the machine are in one-to-one correspondence with the ones of the Crumble GLAM. Finally, we prove that the overall number of the search transitions is bilinear in the number of principal transitions and the size of the initial term, establishing bilinearity of the Pointed Crumble GLAM. To improve the readability of this section, all proofs are moved to Appendix C.5.

Pointed crumbles and pointed environments. The key idea to turn the Crumble GLAM into the Pointed Crumble GLAM is to avoid plugging and unplugging in the rules by letting them act on *pointed crumbles*, i.e. on crumbles where the beginning of the evaluated coda is explicitly marked using a pointer. For example, the crumble (b, ee_v) could be represented as $(b, e \mid e_v)$ where “ \mid ” is the explicit separator that must be followed by v -environments only. A pointed crumble $(b, e[x \leftarrow b'] \mid e_v)$ is the machine state that is attempting to evaluate the crumble $C\langle(b', e_v)\rangle$, where $C = (b, e[x \leftarrow \langle \cdot \rangle])$. If (b', e_v) is a Crumble GLAM a -redex, the Pointed Crumble GLAM will evaluate according to the corresponding a -transition that also takes care of setting (in $O(1)$) the pointer to the rightmost unevaluated crumble. Otherwise, by harmony (Prop. 5.3), b' must be a crumbled

value v and therefore the pointer is moved (in $O(1)$) one step to the left, looking for the next redex: $(b, e[x \leftarrow v] \mid e_v) \rightarrow_{\text{sea}} (b, e \mid [x \leftarrow v]e_v)$.

Not all pointed crumble configurations are of the form $(b, [x \leftarrow b'] \mid e_v)$: the configurations $(b, \mid e_v)$ must be also taken into account and evaluated if b is not a crumbled value. However, there is no simple way to describe machine transitions that act uniformly on configurations $(b, \mid e_v)$ and $(b, e[x \leftarrow b'] \mid e_v)$ without duplicating the rules or without re-introducing a notion of contextual closure. To solve the issue, we abandon pointed crumbles and adopt pointed environments instead.

A pointed environment $([x \leftarrow b]e \mid e_v)$ is just a representation of a pointed crumble $(b, e \mid e_v)$. The leftmost variable x in a pointed environment can be understood as the name given to the machine output. It plays a role similar to the outermost λ -abstraction introduced by CPS transformations, that binds the continuation that is fed with the output of the evaluation. In particular, a normal pointed environment $(\mid [x \leftarrow v]e_v)$ represents the normal crumble (v, e_v) .

Formal definition of pointed environments and read-back. Pointed environments are defined as $e \mid e'$ for any non-pointed environments e and e' such that either e or e' is non-empty. The environment on the left of the cursor \mid is the *unevaluated* environment; the one on the right is the *evaluated* environment.

The translation $\iota(\cdot)$ embeds crumbles into pointed environments:

$$\iota(b, e) := [x \leftarrow b]e \mid \epsilon$$

where x is any variable name fresh in b and e .

The left inverse of $\iota(\cdot)$ is the *read-back* function $(\cdot)_{\parallel}$ (from pointed environments to crumbles):

$$(\epsilon \mid [x \leftarrow b]e)_{\parallel} := (b, e) \quad ([x \leftarrow b]e \mid e')_{\parallel} := (b, ee')$$

Evaluation. The transition rules of the Pointed Crumble GLAM are (the adaptation of) those of the Crumble GLAM plus the new *search* transition \rightarrow_{sea} , and they are all defined only at top level, without a contextual closure—their union is noted \rightarrow_{pCr} :

$$e[x \leftarrow (\lambda y.c)v] \mid e_v \rightarrow_{\beta_v} e[x \leftarrow b]e'[z \leftarrow v] \mid e_v \quad (i)$$

$$e[x \leftarrow \text{if true then } c \text{ else } d] \mid e_v \rightarrow_{\text{ift}} e[x \leftarrow b]e' \mid e_v \quad (ii)$$

$$e[x \leftarrow \text{if false then } c \text{ else } d] \mid e_v \rightarrow_{\text{iff}} e[x \leftarrow b]e' \mid e_v \quad (iii)$$

$$e[x \leftarrow \text{if } v \text{ then } c \text{ else } d] \mid e_v \rightarrow_{\text{ife}} e[x \leftarrow \text{err}] \mid e_v \quad (iv)$$

$$e[x \leftarrow vw] \mid e_v \rightarrow_{@e} e[x \leftarrow \text{err}] \mid e_v \quad (v)$$

$$e[y \leftarrow x] \mid e_v \rightarrow_{\text{subvar}} e[y \leftarrow e_v(x)] \mid e_v \quad (vi)$$

$$e[y \leftarrow xv] \mid e_v \rightarrow_{\text{subl}} e[y \leftarrow e_v(x)v] \mid e_v \quad (vi)$$

$$e[y \leftarrow \text{if } x \text{ then } c \text{ else } d] \mid e_v \rightarrow_{\text{subif}} e[y \leftarrow \text{if } e_v(x) \text{ then } c \text{ else } d] \mid e_v \quad (vii)$$

$$e[x \leftarrow b] \mid e_v \rightarrow_{\text{sea}} e \mid [x \leftarrow b]e_v \quad (vii)$$

where

- $\lambda z.(b, e') := (\lambda y.c)^\alpha$ such that $(e[x \leftarrow b]e'[z \leftarrow v] \mid e_v)$ is well-named.
- where $c := (b, e')$
- where $d := (b, e')$
- if $v = \lambda x.e$ or $v = \text{err}$.
- if $v \in \{\text{true}, \text{false}, \text{err}\}$.

- vi. if $x \in \text{dom}(e_v)$.
- vii. if none of the other rules is applicable, i.e. when b is a practical value $v_{\neg x}$ or when b is x , xv , or if x then c else d but x is not defined in e_v .

A *principal transition* of the Pointed Crumble GLAM is a transition \rightarrow_a for any rule $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e\}$.

PROPOSITION C.12 (HARMONY FOR THE POINTED CRUMBLE GLAM). *A reachable pointed environment e_1 is normal iff it has the form $(e \mid e_v)$ for some non-empty v -environment e_v .*

Implementation. We prove that that the Pointed Crumble GLAM simulates the Crumble GLAM following the standard schema introduced in Subsect. 5.2 and that we have already employed for the Crumble GLAM in Subsect. 5.3. We only state the machine invariants and the statement of the implementation theorem; the details of the proof can be found in the Appendix C.5 (p. 23).

Definition C.13 (Reachable state). A pointed environment is said to be *reachable* if it is obtained from evaluation steps starting from the translation $\iota(c)$ of a well-named closed crumble c .

LEMMA C.14 (INVARIANTS FOR THE POINTED CRUMBLE GLAM). *Let e_1 be a reachable pointed environment in the Pointed Crumble GLAM:*

- (1) Freshness: e_1 is well-named.
- (2) Closure: e_1 is closed.
- (3) Rightmost: $e_1 = (e \mid e_v)$ for some environment e and some v -environment e_v .

THEOREM C.15 (IMPLEMENTATION). *Let e_1 be a reachable pointed environment in the Pointed Crumble GLAM.*

- (1) Initialisation: $(\iota(c))_{\parallel} = c$ for every crumble c .
- (2) Principal projection: if $e_1 \rightarrow_a e'_1$ then $(e_1)_{\parallel} \rightarrow_a (e'_1)_{\parallel}$ for any rule $a \neq \text{sea}$.
- (3) Overhead transparency: if $e_1 \rightarrow_{\text{sea}} e'_1$ then $(e_1)_{\parallel} = (e'_1)_{\parallel}$.
- (4) Determinism: the transition function \rightarrow_{pCr} is deterministic.
- (5) Halt: if e_1 is normal then $(e_1)_{\parallel}$ is normal.
- (6) Overhead termination: \rightarrow_{sea} terminates.

Therefore, the Pointed Crumble GLAM (with its transition function \rightarrow_{pCr}), the Crumble GLAM (with \rightarrow_{Cr}), and the read-back $(\cdot)_{\parallel}$ form an implementation system.

Complexity. We reuse the measures introduced in Subsect. 5.3. The only difference here is that we need to bound also the number of sea steps, which intuitively depends on the length of the pointed environment that is being evaluated. We define the new measure $|\cdot|_{\text{len}}$ on environments and crumbles, that simply counts the number of entries:

$$|b, e|_{\text{len}} := 1 + |e|_{\text{len}} \quad |e|_{\text{len}} := 0 \quad |e[x \leftarrow b]|_{\text{len}} := 1 + |e|_{\text{len}}$$

After each $\beta_v/\text{ift}/\text{iff}$ step, the measure increases by the length of the body that is being concatenated. For every environment e , we define the constant $L(e)$ that bounds the length of the bodies occurring anywhere in e :

$$L(e) := \sup\{|c|_{\text{len}} : c \text{ body in } e\}.$$

We extend the definitions above to pointed environments and crumbles in the expected way:

$$|e \mid e'|_{\text{len}} := |ee'|_{\text{len}} \quad L(e \mid e') := L(ee') \quad L(c) := L(\iota(c)).$$

Remark C.1. For any crumble c and term t : $|c|_{\text{len}} = |\iota(c)|_{\text{len}}$, $L(t) \leq |t|$ and $|t|_{\text{len}} \leq |t|$.

As usual, an execution ρ is a sequence of evaluation steps, and we use $|\rho|_a$ to count the number of a -evaluation steps in ρ .

By Thm. C.15, all transition steps but sea steps are mapped one-to-one to corresponding transition steps in the Crumble GLAM. As for the number of sea transitions:

LEMMA C.16 (NUMBER OF SEA-TRANSITIONS). *Let c be a well-named crumble, and let $\rho: \iota(c) \rightarrow_{\text{pCr}}^* e_1 = (e \mid e_v)$ an execution in the Pointed Crumble GLAM. Then $|e|_{\text{len}} \leq |c|_{\text{len}} + (|\rho|_{\beta_v} + |\rho|_{\text{ift}} + |\rho|_{\text{iff}}) \cdot L(c) - |\rho|_{\text{sea}}$.*

COROLLARY C.17. *Let t be a term. For a normalizing execution ρ in the Pointed Crumble GLAM starting from $\iota(t)$, we have $|\rho|_{\text{sea}} \leq (|\rho|_{\text{p}} + 1) \cdot |t|$.*

Cost of evaluation. The cost of all transitions but sea was already discussed in Subsect. 5.3. The cost of a sea transition is clearly $O(1)$.

THEOREM C.18 (THE POINTED CRUMBLE GLAM IS BILINEAR). *For any closed term t and any Pointed Crumble GLAM execution $\rho: \iota(t) \rightarrow_{\text{pCr}}^* e_1$, the cost of implementing ρ on a RAM is $O((|\rho|_{\text{p}} + 1) \cdot |t|)$.*

The previous theorem also has a consequence for the non-pointed case. Since the theorem shows that the cost of searching for redexes and the (un)plugging operations can indeed be realised in bilinear time, it is also true that the Crumble GLAM can be implemented in bilinear time, improving Thm. 5.10 by removing the up to search and (un)plugging side condition.

COROLLARY C.19 (THE CRUMBLE GLAM IS BILINEAR). *For any closed term t and any Crumble GLAM execution $\rho: t \rightarrow_{\text{Cr}}^* c$, the cost of implementing ρ on a RAM is $O((|\rho|_{\text{p}} + 1) \cdot |t|)$.*

C.5 Proofs of Appendix C.4 (p. 22)

LEMMA C.20. *Let b be a bite and e an environment. $[x \leftarrow b]e$ is an v -environment if and only if (b, e) is a v -crumble.*

PROOF. Obvious from the definition of v -environment and v -crumble. \square

PROOF OF (LEMMA C.14). *Let e_1 be a reachable pointed environment in the Pointed Crumble GLAM:*

- (1) Freshness: e_1 is well-named.
- (2) Closure: e_1 is closed.
- (3) Rightmost: $e_1 = (e \mid e_v)$ for some environment e and some v -environment e_v .

PROOF. By induction on the length of the evaluation sequence leading to e_1 . The base cases hold by the definition of reachability and by the definition of $\iota(\cdot)$. As for the inductive cases, we proceed by cases on the transition rules:

- (1) For β_v the claim follows from the side condition. The rules in $\{\text{ift}, \text{iff}, \text{ife}, @e, \text{sea}\}$ do not increase the number of explicit substitutions occurring in e_1 outside of abstractions, hence the claim follows from the i.h.. The rules in $\{\text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$ copy a practical value from the environment e_v : note that

$e_v(x)$ is either an abstraction (which does not influence well-namedness) or a boolean or an error (which do not contain explicit substitutions).

- (2) Similar to the discussion in Lemma 5.5.
- (3) The rules in $\{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e, \text{sub}_{var}, \text{sub}_I, \text{sub}_{if}\}$ do not change the evaluated part, hence the claim follows from the *i.h.*. As for the sea rule, by Lemma C.20 it suffices to prove that (b, e_v) is a v -crumble, knowing that the other transition rules cannot be applied. This follows from the side condition of the sea rule: the crumble b is necessarily a practical value, as the other cases in which the free variable x occurs are not possible because e_1 is a closed pointed environment. \square

PROOF OF (PROP. C.12). A reachable pointed environment e_1 is normal iff it has the form $(\epsilon \mid e_v)$ for some non-empty v -environment e_v .

PROOF. The proof of the implication from right to left is trivial. Let us now prove the other direction. Let e_1 be a reachable normal pointed environment. By Lemma C.14, e_1 has the form $(\epsilon \mid e_v)$. e cannot be non-empty, because otherwise one of the transition rules in $\{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e, \text{sub}_{var}, \text{sub}_I, \text{sub}_{if}, \text{sea}\}$ could be applied, contradicting the hypothesis that e_1 is normal. \square

LEMMA C.21. For every pointed environment $(e[x \leftarrow b] \mid e')$: $(e[x \leftarrow b] \mid e')_{\Downarrow} = (C(b, e'))_{\Downarrow}$ where C is:

$$C := \begin{cases} \langle \cdot \rangle & \text{if } e = \epsilon \\ (b', e''[x \leftarrow \langle \cdot \rangle]) & \text{if } e := [y \leftarrow b']e'' \end{cases}$$

PROOF. Easy by the definition of $(\cdot)_{\Downarrow}$. \square

PROOF OF (THM. C.15). Let e_1 be a reachable pointed environment in the Pointed Crumble GLAM.

- (1) Initialisation: $(\iota(c))_{\Downarrow} = c$ for every crumble c .
- (2) Principal projection: if $e_1 \rightarrow_a e'_1$ then $(e_1)_{\Downarrow} \rightarrow_a (e'_1)_{\Downarrow}$ for any rule $a \neq \text{sea}$.
- (3) Overhead transparency: if $e_1 \rightarrow_{\text{sea}} e'_1$ then $(e_1)_{\Downarrow} = (e'_1)_{\Downarrow}$.
- (4) Determinism: the transition function \rightarrow_{pCr} is deterministic.
- (5) Halt: if e_1 is normal then $(e_1)_{\Downarrow}$ is normal.
- (6) Overhead termination: \rightarrow_{sea} terminates.

Therefore, the Pointed Crumble GLAM (with its transition function \rightarrow_{pCr}), the Crumble GLAM (with \rightarrow_{Cr}), and the read-back $(\cdot)_{\Downarrow}$ form an implementation system.

PROOF.

- (1) Let $c = (b, e)$: by the definitions, $(\iota(c))_{\Downarrow} = ([x \leftarrow b]e \mid \epsilon)_{\Downarrow} = (b, e\epsilon) = c$.
- (2) There is a clear one-to-one correspondence between the transitions of the Pointed Crumble GLAM and the Crumble GLAM (apart from sea). In order to prove it, one may proceed by cases on each transition rule, but we just show the case of β_v as the others are similar. Suppose $e[x \leftarrow (\lambda y.c)v] \mid e_v \rightarrow_{\beta_v} e[x \leftarrow b]e'[z \leftarrow v] \mid e_v$. By Lemma C.21, $(e[x \leftarrow (\lambda y.c)v] \mid e_v)_{\Downarrow} = C(((\lambda y.c)v, e_v))$ and $(e[x \leftarrow b]e'[z \leftarrow v] \mid e_v)_{\Downarrow} = C(b, e'[z \leftarrow v]e_v)$ for some crumble context C . Clearly also $C(((\lambda y.c)v, e_v)) \rightarrow_{\beta_v} C(b, e'[z \leftarrow v]e_v)$.

- (3) By inspection of the rule sea. We need to prove that $(e[x \leftarrow b] \mid e_v)_{\Downarrow} = (e \mid [x \leftarrow b]e_v)_{\Downarrow}$. By cases on the structure of e : if $e = \epsilon$, then $(\epsilon[x \leftarrow b] \mid e_v)_{\Downarrow} = (b, e_v) = (\epsilon \mid [x \leftarrow b]e_v)_{\Downarrow}$. If instead $e = [y \leftarrow b']e'$, then $(e[x \leftarrow b] \mid e_v)_{\Downarrow} = (b', e'[x \leftarrow b]e_v) = (e \mid [x \leftarrow b]e_v)_{\Downarrow}$.
- (4) The rule sea can be applied by definition only when the other rules cannot be applied. The remaining rules apply to a pointed environment of the form $(e[x \leftarrow b] \mid e_v)$, for distinct shapes of b :
 - The rules $\{\beta_v, @e, \text{sub}_I\}$ apply when b is an application, and respectively the application of an abstraction to a crumbled value, of a boolean/error to a crumbled value, and of a variable to a crumbled value.
 - The rule sub_{var} applies when b is a variable.
 - The rules in $\{\text{ift}, \text{iff}, \text{ife}, \text{sub}_{if}\}$ apply when b is a if-then-else, and in clearly disjoint cases according to the structure of the condition: respectively, when the condition is true, false, an abstraction or an error, and a variable.
 - Finally, note that no single rule in $\{\text{sub}_{var}, \text{sub}_I, \text{sub}_{if}\}$ can transition to different pointed environments due to the lookup of a variable which has multiple occurrences in the environment: the lookup is deterministic because the environment is well-named (Lemma C.14.1).
- (5) By Prop. C.12, e_1 is normal iff it has the form $(\epsilon \mid e_v)$ for some non-empty v -environment e_v . Then $(e_1)_{\Downarrow} = (\epsilon \mid e_v)_{\Downarrow}$ which is a v -crumble by Lemma C.20. By Prop. 5.3, $(e_1)_{\Downarrow}$ is normal.
- (6) Immediate consequence of forthcoming Cor. C.17 (proved independently). \square

PROOF OF (LEMMA C.16). Let c be a well-named crumble, and let $\rho: \iota(c) \rightarrow_{\text{pCr}}^* e_1 = (\epsilon \mid e_v)$ an execution in the Pointed Crumble GLAM. Then $|e|_{\text{len}} \leq |c|_{\text{len}} + (|\rho|_{\beta_v} + |\rho|_{\text{ift}} + |\rho|_{\text{iff}}) \cdot L(c) - |\rho|_{\text{sea}}$.

PROOF. By induction on the length of ρ . In the base case $|e|_{\text{len}} = |c|_{\text{len}}$ and $|\rho|_{\beta_v} = |\rho|_{\text{ift}} = |\rho|_{\text{iff}} = |\rho|_{\text{sea}} = 0$. In the inductive case, use the *i.h.* and proceed by cases on the transition rules. The transitions $\text{sub}_{var}, \text{sub}_I, \text{sub}_{if}$ do not change the measure of the unevaluated environment because $e_v(x)$ is a value, whose measure is 0 by definition. A sea transition decreases by 1 the measure of the unevaluated part. A β_v/ift transition increases the measure of the unevaluated part by a number bound by $L(c)$. A iff transition increases the measure of the unevaluated part by a number bound by $L(d)$. All remaining transitions do not increase the measure of the unevaluated environment. \square

PROOF OF (THM. C.18). For any closed term t and any Pointed Crumble GLAM execution $\rho: \iota(t) \rightarrow_{\text{pCr}}^* e_1$, the cost of implementing ρ on a RAM is $O((|\rho|_p + 1) \cdot |t|)$.

PROOF. The cost $|\rho|$ of ρ is the the total cost of principal transitions (which was proved in Subsect. 5.3 to be bilinear in the number of principal steps (plus one) and the size of the initial crumble when starting from terms) plus the total cost of sea transitions. The cost of sea transitions in ρ is linear in the number of sea transitions, and therefore by Cor. C.17, again bilinear in the number of principal steps (plus one) and the size of the initial term. \square

D PROOFS OF SECT. 6 (OPEN CASE)

D.1 Proofs of Subsect. 6.1

LEMMA D.1 (COMPOSITION OF RIGHT F-CONTEXTS). *Let R and R' be right f -contexts. Then their composition $R\langle R' \rangle$ is a right f -context.*

PROOF. By induction on the right f -context R . Cases:

- *Hole*, i.e. $R := \langle \cdot \rangle$: then, $R\langle R' \rangle = R'$ is a right f -context by hypothesis.
- *Right*, i.e. $R := tR''$: then, $R\langle R' \rangle = t(R''\langle R' \rangle)$ is a right f -context because $R''\langle R' \rangle$ is a right f -context by i.h.
- *Left*, i.e. $R := R''f$: then, $R\langle R' \rangle = R''\langle R' \rangle f$ is a right f -context because $R''\langle R' \rangle$ is a right f -context by i.h.
- *if-then-else*, i.e. $R := (\text{if } R'' \text{ then } t \text{ else } u)$: then, $R\langle R' \rangle = (\text{if } R''\langle R' \rangle \text{ then } t \text{ else } u)$ is a right f -context because $R''\langle R' \rangle$ is a right f -context by i.h. \square

PROOF OF (PROP. 6.2). *Let t, u be terms.*

- (1) Open harmony: t is $c\beta_f$ -normal if and only if t is a fireball.
- (2) Inert substitutions and evaluation commute: *Let i be an inert term. Then $t \rightarrow_{c\beta_f} u$ if and only if $t\{x \leftarrow i\} \rightarrow_{c\beta_f} u\{x \leftarrow i\}$.*

PROOF. (1) We show separately the two implications:

(\Rightarrow) Proof by induction on the structure of t .

- Case t value: trivial.
- Case $t = us$ for some terms u and s : we show that then t is an inert term (and hence a fireball). Since t is $c\beta_f$ -normal, then s is $c\beta_f$ -normal, and hence by i.h. s is a fireball. Since t is $c\beta_f$ -normal and s is a fireball, then also u must be $c\beta_f$ -normal, and hence a fireball by i.h. We proceed by cases on u , showing that the only possible cases are when u is a variable or an inert term (and so t is a fireball). Indeed, u cannot be an abstraction (because otherwise the rule \rightarrow_{β_f} may be applied, contradicting the hypothesis that t is $c\beta_f$ -normal); it cannot be a boolean or err (because otherwise the rule $\rightarrow_{@e}$ may be applied).
- Case $t = \text{if } u \text{ then } s \text{ else } r$ for some terms u, s, r : we show that the only possible cases are when u is a variable or an inert term (and hence t is a fireball). Since t is $c\beta_f$ -normal, then u is $c\beta_f$ -normal, and hence by i.h. u is a fireball. u it cannot be an abstraction or err (because otherwise the rule \rightarrow_{ife} may be applied, contradicting the hypothesis that t is $c\beta_f$ -normal); it cannot be a boolean (because otherwise one of the rules \rightarrow_{ift} or \rightarrow_{iff} may be applied).

(\Leftarrow) By hypothesis, t is a fireball, i.e. either a value or an inert term.

- *Value*: We proceed by cases on the definition of value. If t is a variable, then it is clearly $c\beta_f$ -normal. If t is an abstraction, then it is $c\beta_f$ -normal since $\rightarrow_{c\beta_f}$ does not reduce under λ 's. Otherwise t is a either a boolean or err, which clearly are $c\beta_f$ -normal.
- *Inert term*: We proceed by cases on the definition of inert term. If $t = xf$ for some variable x and some fireball f , then f is $c\beta_f$ -normal by i.h., and x is clearly $c\beta_f$ -normal and different from abstractions, booleans and err; therefore, no reduction rule can be applied to

t , i.e. t is $c\beta_f$ -normal. If $t = xf$ or $t = if$ for some inert term i and some fireball f , then i and f are $c\beta_f$ -normal by i.h.; moreover, i is not an abstraction or boolean or err; therefore, no reduction rule can be applied to t , i.e. t is $c\beta_f$ -normal. Otherwise, $t = \text{if } u \text{ then } s \text{ else } r$ for some terms u, s, r where u is either a variable or a inert term; since u is $c\beta_f$ -normal (if u is a inert term, this holds by i.h.), then no reduction rule can be applied to t , i.e. t is $c\beta_f$ -normal. \square

- (2) The direction (\Leftarrow) is exactly Lemma D.3. The direction (\Rightarrow) is proved by induction on the definition of $t \rightarrow_{c\beta_f} u$, in a way that is analogue to the proof of Lemma D.3.

LEMMA D.2 (FIREBALLS ARE CLOSED UNDER SUBSTITUTION AND ANTI-SUBSTITUTION OF INERT TERMS). *Let t be a term and i be an inert term. Then, $t\{x \leftarrow i\}$ is a fireball (resp. an abstraction) if and only if t is a fireball (resp. an abstraction). Moreover, $t\{x \leftarrow i\} = \text{true}$ (resp. $t\{x \leftarrow i\} = \text{false}$; $t\{x \leftarrow i\} = \text{err}$) if and only if $t = \text{true}$ (resp. $t = \text{false}$; $t = \text{err}$).*

PROOF. The left-to-right direction (\Rightarrow) is proved by a simple induction on the fireball structure of $t\{x \leftarrow i\}$. Conversely, the right-to-left direction (\Leftarrow) is proved by a simple induction on the inert structure of t . \square

LEMMA D.3 (SUBSTITUTION OF INERT TERMS DOES NOT CREATE REDEXES). *Let t, u be terms and i be an inert term. There is a term s such that, if $t\{x \leftarrow i\} \rightarrow_{c\beta_f} u$ then $t \rightarrow_{c\beta_f} s$ and $u = s\{x \leftarrow i\}$.*

PROOF. By induction on the right fireball context closing the redex. Cases:

- *Step at the root*: Sub-cases:

- (1) *Beta-step*, i.e. $t\{x \leftarrow i\} = (\lambda y. r\{x \leftarrow i\})q\{x \leftarrow i\} \rightarrow_{\beta_f} r\{x \leftarrow i\}\{y \leftarrow q\{x \leftarrow i\}\} =: u$ where $t := (\lambda x. r)q$ and $q\{x \leftarrow i\}$ is a fireball. By Lemma D.2, q is a fireball and hence $t = (\lambda x. r)q \rightarrow_{c\beta_f} r\{y \leftarrow q\}$. So, $s := r\{y \leftarrow q\}$ satisfies the statement, as $s\{x \leftarrow i\} = r\{y \leftarrow q\}\{x \leftarrow i\} = u$.
- (2) *Conditional step with true*, i.e. $t\{x \leftarrow i\} = \text{if true then } s\{x \leftarrow i\} \text{ else } q\{x \leftarrow i\} \rightarrow_{ift} s\{x \leftarrow i\} =: u$ where $t := \text{if } r \text{ then } s \text{ else } q$ for some term r such that $r\{x \leftarrow i\} = \text{true}$. By Lemma D.2, $r = \text{true}$ and hence $t = \text{if true then } s \text{ else } q \rightarrow_{ift} s$.
- (3) *Conditional step with false*, i.e. $t\{x \leftarrow i\} = \text{if true then } q\{x \leftarrow i\} \text{ else } s\{x \leftarrow i\} \rightarrow_{iff} s\{x \leftarrow i\} =: u$ where $t := \text{if } r \text{ then } q \text{ else } s$ for some term r such that $r\{x \leftarrow i\} = \text{false}$. By Lemma D.2, $r = \text{false}$ and hence $t = \text{if false then } q \text{ else } s \rightarrow_{iff} s$.
- (4) *Conditional step to error*, i.e. $t\{x \leftarrow i\} = \text{if } r \text{ then } q\{x \leftarrow i\} \text{ else } m\{x \leftarrow i\} \rightarrow_{ife} \text{err} =: u$ where $r := \lambda y. r'$ or $r = \text{err}$, and $t := \text{if } p \text{ then } q \text{ else } m$ for some term p such that $p\{x \leftarrow i\} = r$. By Lemma D.2, p is either an abstraction or err and hence $t = \text{if } r \text{ then } q \text{ else } m \rightarrow_{ife} \text{err}$. So, $s := \text{err}$ satisfies the statement, as $s\{x \leftarrow i\} = \text{err} = u$.
- (5) *Application step to error*, i.e. $t\{x \leftarrow i\} = r\{x \leftarrow i\}q\{x \leftarrow i\} \rightarrow_{@e} \text{err} =: u$ where $r\{x \leftarrow i\} \in \{\text{true}, \text{false}, \text{err}\}$ and $t := rq$. By Lemma D.2, $r \in \{\text{true}, \text{false}, \text{err}\}$ and hence $t = rq \rightarrow_{@e} \text{err}$. So, $s := \text{err}$ satisfies the statement, as $s\{x \leftarrow i\} = \text{err} = u$.

- *Application left*, i.e. $t\{x \leftarrow i\} = r\{x \leftarrow i\}f\{x \leftarrow i\} \rightarrow_{c\beta_f} pf\{x \leftarrow i\} =: u$ where $t := rf$ and $r\{x \leftarrow i\} \rightarrow_{c\beta_f} p$. By i.h. there is a term s' such that $p = s'\{x \leftarrow i\}$ and $r \rightarrow_{c\beta_f} s'$. Then $s := s'f$ satisfies the statement, as $s\{x \leftarrow i\} = s'\{x \leftarrow i\}fxi = u$.
- *Application right*, i.e. $t\{x \leftarrow i\} = q\{x \leftarrow i\}r\{x \leftarrow i\} \rightarrow_{c\beta_f} q\{x \leftarrow i\}p =: u$ where $t := qr$ and $r\{x \leftarrow i\} \rightarrow_{c\beta_f} p$. Analogous to the application left case, just switch left and right.
- *If-then-else*, i.e. $t\{x \leftarrow i\} = \text{if } r\{x \leftarrow i\} \text{ then } q\{x \leftarrow i\} \text{ else } p\{x \leftarrow i\} \rightarrow_{c\beta_f} \text{if } m \text{ then } q\{x \leftarrow i\} \text{ else } p\{x \leftarrow i\} =: u$ where $t := \text{if } r \text{ then } q \text{ else } p$ and $r\{x \leftarrow i\} \rightarrow_{c\beta_f} m$. By i.h. there is a term s' such that $m = s'\{x \leftarrow i\}$ and $r \rightarrow_{c\beta_f} s'$. Then $s := \text{if } s' \text{ then } q \text{ else } p$ satisfies the statement, as $s\{x \leftarrow i\} = \text{if } s'\{x \leftarrow i\} \text{ then } q\{x \leftarrow i\} \text{ else } p\{x \leftarrow i\} = u$. \square

D.2 Proofs of Subsect. 6.2

D.2.1 Evaluated environments.

Given the terms t_1, \dots, t_n and pairwise distinct variables x_1, \dots, x_n , let $\sigma := \{x_1 \leftarrow t_1, \dots, x_n \leftarrow t_n\}$ be the *simultaneous* capture-avoiding substitution of t_i for x_i , for all $1 \leq i \leq n$. In particular, if $n = 0$ then $\sigma = \{\}$, the identity (or empty) substitution. If u is a term and $y \neq x_i$ for all $1 \leq i \leq n$, we set $\{y \leftarrow u\} \cup \sigma := \{y \leftarrow u, x_1 \leftarrow t_1, \dots, x_n \leftarrow t_n\}$.

Given an environment e , the simultaneous substitution $\sigma(e)$ associated with e is defined by induction on the length of e :

$$\sigma(\epsilon) := \{\} \quad \sigma([x \leftarrow b]e) := \{x \leftarrow (b, e)_\downarrow\} \cup \sigma(e).$$

LEMMA D.4 (SEMI-CLOSURE UNDER SUBSTITUTION). *Let v be a value, x_1, \dots, x_n be pairwise distinct variables, and f_1, \dots, f_n be fireballs. Then $v\{x_1 \leftarrow f_1, \dots, x_n \leftarrow f_n\}$ is a fireball. If moreover v is a λ -abstraction, then $v\{x_1 \leftarrow f_1, \dots, x_n \leftarrow f_n\}$ is a λ -abstraction.*

PROOF. Let $\sigma := \{x_1 \leftarrow f_1, \dots, x_n \leftarrow f_n\}$. By cases on the definition of value:

- *Variable*, i.e. either $v := x_i$ for some $1 \leq i \leq n$ and then $v\sigma = f_i$ which is a fireball by hypothesis; or $v := y \neq x_i$ for all $1 \leq i \leq n$ and then $v\sigma = y$ which is a fireball.
- *Abstraction*, i.e. $v := \lambda y. u$ and we can suppose without loss of generality that $y \notin \bigcup_{i=1}^n \text{fv}(f_i) \cup \{x_1, \dots, x_n\}$; therefore, $v\sigma = \lambda y. (u\sigma)$ which is a λ -abstraction and hence a fireball.
- *Booleans and errors*, i.e. $v \in \{\text{true}, \text{false}, \text{err}\}$. Trivial since $v\sigma = v$. \square

LEMMA D.5 (READ-BACK TO FIREBALL). *For every crumbled value v and f -environment e_f , one has that $(v, e_f)_\downarrow$ is a fireball. If moreover v is a practical value, then $(v, e_f)_\downarrow$ is a practical value.*

PROOF. By induction on the length of e_f . We proceed by cases on the shape of v .

- *Abstraction*: If $e_f := \epsilon$, then clearly $(v, e_f)_\downarrow = v_\downarrow$ is a λ -abstraction and hence a fireball. Otherwise $e_f := e'_f[x \leftarrow b]$ where b is a bite such that b_\downarrow is a fireball: thus, $(v, e_f)_\downarrow = (v, e'_f)_\downarrow\{x \leftarrow b_\downarrow\}$; by i.h., $(v, e'_f)_\downarrow$ is a λ -abstraction, thus $(v, e_f)_\downarrow$ is a λ -abstraction (and so a fireball) by Lemma D.4.

- *Booleans and errors*: the proof is identical to the previous case.
- *Variable*: If $v := x \notin \text{dom}(e_f)$, then $(v, e_f)_\downarrow = x$ which is a fireball; otherwise $v := x \in \text{dom}(e_f)$ with $e_f := e'_f[x \leftarrow b]e''_f$, and then $(x, e_f)_\downarrow = (b, e''_f)_\downarrow$ (since $x \notin \text{dom}(e'_f)$) is a fireball by definition of f -environment, as the f -environment $[x \leftarrow b]e''_f = E\langle(b, e''_f)\rangle$ with $E := [x \leftarrow \langle \cdot \rangle]$. \square

LEMMA D.6. *Let e_f well-named. If $(x, e_f)_\downarrow$ is a practical value, then $e_f(x)$ is a practical value.*

PROOF. Assume by contradiction that x is not defined in e_f : then by Lemma B.5 $(x, e_f)_\downarrow = (x, \epsilon)_\downarrow = x$, contradicting the hypothesis that $(x, e_f)_\downarrow$ is a practical value. Therefore x must be defined in e_f , i.e. $e_f := e'_f[x \leftarrow b]e''_f$ with $b := e_f(x)$. By the hypothesis that e_f is well-named, $x \notin \text{dom}(e'_f)$; therefore $(x, e_f)_\downarrow = (b, e''_f)_\downarrow$ by Lemma B.9. By the definition of e_f , since $(b, e''_f)_\downarrow$ is a practical value, then also b is a practical value, and we conclude. \square

LEMMA D.7. *Let $c = (b, e_f)$ be a well-named crumble, and let b have the following property: either b is a practical value, or b is x or xv or if x then c else d but x is not defined in e_f or $e_f(x)$ is not a practical value. Then c is a f -crumble.*

PROOF. Let $c = (b, e_f)$ as above: it suffices to prove that $(b, e_f)_\downarrow$ is a fireball, and that if $(b, e_f)_\downarrow$ is a practical value, then also b is a practical value. By cases on the property about b in the hypothesis:

- *Abstraction*: nothing to prove because $(b, e_f)_\downarrow$ is an abstraction by Lemma D.5 and thus a fireball.
- *Booleans and errors*: nothing to prove because $(b, e_f)_\downarrow$ is a boolean or an error and thus a fireball.
- *Variable*, i.e. $b = x$ for some x when x is not defined in e_f or $e_f(x)$ is not a practical value. $(b, e_f)_\downarrow$ is a fireball by Lemma D.5. Let us now assume that $(b, e_f)_\downarrow$ is a practical value, and show that it is not possible: in fact by Lemma D.6 $e_f(x)$ must then be defined and a practical value, contradicting the hypothesis.
- *Application of a variable to a crumbled value*, i.e. $b = xv$ when x is not defined in e_f or $e_f(x)$ is not a practical value. Note that $(b, e_f)_\downarrow = (x, e_f)_\downarrow(v, e_f)_\downarrow$, where both $(x, e_f)_\downarrow$ and $(v, e_f)_\downarrow$ are fireballs by Lemma D.5. If $(x, e_f)_\downarrow$ is inert there is nothing else to prove, because then $(b, e_f)_\downarrow$ is a fireball, and clearly not a practical value. The case when $(x, e_f)_\downarrow$ is a practical value is not possible: again by Lemma D.6 $e_f(x)$ should be defined and a practical value, contradicting the hypothesis.
- *if-then-else* where the guard is a variable, i.e. $b = \text{if } x \text{ then } c \text{ else } d$ when x is not defined in e_f or $e_f(x)$ is not a practical value. Note that $(b, e_f)_\downarrow = \text{if } (x, e_f)_\downarrow \text{ then } (c @ e_f)_\downarrow \text{ else } (d @ e_f)_\downarrow$, where $(x, e_f)_\downarrow$ is a fireball by Lemma D.5. If $(x, e_f)_\downarrow$ is inert there is nothing else to prove, because then $(b, e_f)_\downarrow$ is a fireball, and clearly not a practical value. The case when $(x, e_f)_\downarrow$ is a practical value is not possible: again by Lemma D.6 $e_f(x)$ should be defined and a practical value, contradicting the hypothesis. \square

COROLLARY D.8. *If the well-named crumble (b, e_f) is normal, then it is a f -crumble.*

PROOF OF (PROP. 6.3). *A crumble c is oCr-normal if and only if it is a f -crumble.*

PROOF.

- (\Rightarrow) Let $c = (b, e)$ be well-named and normal. We proceed by structural induction on e :
- if $e = \epsilon$, then (b, ϵ) is a f -crumble by Cor. D.8;
 - if $e = [x \leftarrow b']e'$, then also the crumble (b', e') is normal. By *i.h.* (b', e') is a f -crumble, and therefore $e = [x \leftarrow b']e'$ is a f -environment. By Cor. D.8, $c = (b, e)$ is a f -crumble.
- (\Leftarrow) Let $c = c_v$, we need to prove that c_v is normal. Let $c_v = C\langle(b, e_f)\rangle$ for some C, b, e_f . First of all, note that by the definition of c_v , $(b, e_f)_\downarrow$ is a fireball, and that if $(b, e_f)_\downarrow$ is a practical value, then also b is a practical value. We prove that no reduction rule is applicable to $C\langle(b, e_f)\rangle$:
- Rule β_f can be applied only if $b = (\lambda x.b')v$, but this contradicts the hypothesis that $(b, e_f)_\downarrow$ is a fireball, since $((\lambda x.b')v, e_f)_\downarrow = (\lambda x.b', e_f)_\downarrow(v, e_f)_\downarrow$ and $(\lambda x.b', e_f)_\downarrow$ is an abstraction by Lemma D.5.
 - The analysis for the rules *ift*, *iff*, *ife*, *@e* is similar to the previous case: in all cases the read-back of the lhs of the rule is not a fireball, contradicting the hypothesis.
 - Rule sub_{var} can be applied only if b is some variable x and $e_f(x)$ is defined and a practical value. This contradicts the hypothesis that if $(b, e_f)_\downarrow$ is a practical value, then also b is a practical value.
 - Rule sub_l can be applied only if $b = xv$ for some x, v , and $e_f(x)$ is defined and a practical value. This contradicts the hypothesis that $(b, e_f)_\downarrow$ is a fireball, since $(xv, e_f)_\downarrow = (x, e_f)_\downarrow(v, e_f)_\downarrow$ and $(x, e_f)_\downarrow$ is a practical value by Lemma D.5 because $e_f(x)$ is a practical value.
 - Rule sub_{if} can be applied only if $b = \text{if } x \text{ then } c \text{ else } d$ for some x, c, d , and $e_f(x)$ is defined and a practical value. This contradicts the hypothesis that $(b, e_f)_\downarrow$ is a fireball, since $(\text{if } x \text{ then } c \text{ else } d, e_f)_\downarrow = \text{if } (x, e_f)_\downarrow \text{ then } (c @ e_f)_\downarrow \text{ else } (d @ e_f)_\downarrow$ and $(x, e_f)_\downarrow$ is a practical value by Lemma D.5 because $e_f(x)$ is a practical value. \square

D.2.2 The implementation theorem for the Open Crumble GLAM.

LEMMA D.9 (INVARIANTS FOR THE OPEN CRUMBLE GLAM). *For every reachable crumble c :*

- (1) Freshness: c is well-named.
- (2) Disjointedness: if $c = C\langle(b, e)\rangle$ then $\text{dom}(C) \cap \text{fv}(b) = \emptyset$.
- (3) Bodies: any body in c is a subterm (up to renaming) of the initial crumble.
- (4) Weak contextual decoding: for every decomposition $C\langle(b, e_f)\rangle$ where $(b, e_f)_\downarrow$ is not a fireball, if C'' is a prefix of C then C''_\downarrow is a right v-context.

PROOF. By induction on the length of the reduction sequence leading to the crumble c . The base cases hold by Lemma 4.5 (by noting that for Point 4, Lemma 4.5.5 implies the weaker statement Lemma D.9.4). As for the inductive cases, we inspect each reduction rule:

1. The substitution transitions sub_{var} , sub_l , sub_{if} do not change the set of variables occurring on the lhs of substitutions outside abstractions because they copy a value that does not contain any. Hence the claim follows from the *i.h.*. For the rule β_f the claim follows from the side condition. For the remaining rules *ift*, *iff*, *ife*, *@e* the claim follows from the fact that all substitutions outside abstractions in the rhs already occur in the lhs.
2. The substitution transitions sub_{var} , sub_l , sub_{if} do not change the domain of the crumble and only copy to the left a value v such that $c = C'\langle(v, e')\rangle$ and C is a prefix of C' . Thus the claim follows from the *i.h.* because $\text{dom}(C) \subseteq \text{dom}(C')$ and the free variables of the rhs are a subset of the union of the free variables of the lhs with $\text{fv}(v)$.
Transition β_f copies to the toplevel and renames the body of an abstraction. By the properties of α -renaming $\text{fv}((c @ [x \leftarrow v])^\alpha) = \text{fv}(c @ [x \leftarrow v]) = \text{fv}(\lambda x.c)$. If the b is chosen to be in the crumble context (say C'') or in e_f of the reduction rule, then the claim follows from the *i.h.*. Let instead $(c @ [x \leftarrow v])^\alpha =: C'\langle(b, e')\rangle$ with $C = C''\langle C'\rangle$: then $\text{fv}(C'\langle(b, e')\rangle) = \text{fv}(\lambda x.c)$ and $\text{fv}(b) \subseteq \text{dom}(e') \cup \text{fv}(\lambda x.c)$. $\text{dom}(e') \cap \text{dom}(C) = \emptyset$ by the side condition of β_f , and $\text{fv}(\lambda x.c) \cap \text{dom}(C) = \emptyset$ by *i.h.*, therefore we conclude with $\text{fv}(b) \cap \text{dom}(C) = \emptyset$.
3. Transitions sub_{var} , sub_l , sub_{if} may copy an abstraction, but the abstraction was already in the environment, and the claim follows from the *i.h.*. Transition β_f copies and renames the body of an abstraction that was already in the environment, and the claim follows from the *i.h.* since the translation commutes with the renaming of free variables (Remark 4.1.3). All the bodies in the rhs of the remaining transitions *ift*, *iff*, *ife*, *@e* already occur in the lhs and therefore the claim follows from the *i.h.*.
4. Let $b'' \rightarrow^n C'\langle(b', e'_v)\rangle \rightarrow_a C\langle(b, e_v)\rangle$ (where b is not a practical value). Cases of the transition $C'\langle(b', e'_v)\rangle \rightarrow_a C\langle(b, e_v)\rangle$:
 - Case β_f : $C'\langle((\lambda x.c)v, e_v)\rangle \rightarrow_{\beta_v} C'\langle c^\alpha @ ([x \leftarrow v]e_v)\rangle$.
Let C'' be a prefix of C . There are two sub-cases:
 - * C'' is a prefix of C' : by *i.h.* C''_\downarrow is a right v-context.
 - * C' is a prefix of C'' , i.e. $C'' = C'\langle C''' \rangle$ and $c^\alpha = C''' \langle c' \rangle$.
By Lemma 4.5.4 and Lemma D.9.3 c is the translation of a λ -term, by Remark 4.1.3 c^α is so, and thus by Lemma 4.5.5 C'''_\downarrow is a right v-context. By *i.h.*, C'_\downarrow is a right v-context as well. Since $C''_\downarrow = C'_\downarrow \langle C'''_\downarrow \rangle$ according to Corollary B.11.2, we obtain that C''_\downarrow is a right v-context as composition of right v-contexts (Lemma A.1).
 - Case *ift*: $C'\langle(\text{if true then } c \text{ else } d, e_v)\rangle \rightarrow_{\beta_v} C'\langle c @ e_v \rangle$.
Let C'' be a prefix of C . There are two sub-cases:
 - * C'' is a prefix of C' : by *i.h.* C''_\downarrow is a right v-context.
 - * C' is a prefix of C'' , i.e. $C'' = C'\langle C''' \rangle$ and $c = C''' \langle c' \rangle$.
By Lemma 4.5.4 and Lemma D.9.3 c is the translation of a λ -term and thus by Lemma 4.5.5 C'''_\downarrow is a right v-context. By *i.h.*, C'_\downarrow is a right v-context as well. Since $C''_\downarrow = C'_\downarrow \langle C'''_\downarrow \rangle$ according to Corollary B.11.2, we obtain

that C''_{\downarrow} is a right v -context as composition of right v -contexts (Lemma A.1).

- Case iff : identical to the previous case.
- Cases sub_{var} , ife , @e : they follow from the *i.h.* since C is necessarily a prefix of C' because b is a practical value.
- Cases sub_l and sub_{if} : they follow from the *i.h.*, since $e'_v = e_v$ and $C = C'$. \square

LEMMA D.10 (DETERMINISM). \rightarrow_{ocr} is deterministic.

PROOF. Assume that there exists a crumble that may be decomposed in two ways $C\langle(b, e_f)\rangle = C'\langle(b', e'_f)\rangle$ such that they reduce respectively $C\langle(b, e_f)\rangle \rightarrow_a C\langle c \rangle$ and $C'\langle(b', e'_f)\rangle \rightarrow_b C'\langle d \rangle$ with rules $a, b \in \{\beta_f, \text{ift}, \text{iff}, \text{ife}, \text{@e}, \text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$.

We prove that it must necessarily be $a = b$, $C = C'$, and $c = d$ (up to α -equivalence). Three cases:

- C strict initial segment of C' , i.e. $C' = C\langle C'' \rangle$ for some $C'' \neq \langle \cdot \rangle$. We show that this case is not possible: in fact, it follows that $e_f = E\langle(b', e'_f)\rangle$ for some E , thus (b', e'_f) is a f -crumble, and by Prop. 6.3 it must be normal, contradicting the hypothesis that (b', e'_f) and c reduce with rule b .
- $C = C'$. By inspection of the reduction rules, $a = b$: in fact the rule β_f applies only when b is the application of an abstraction to a crumbled value, the rule sub_{var} only when b is a variable, and the rule sub_l only when b is the application of a variable to a crumbled value, etc. It remains to show that $c = d$ (up to α): this follows from the determinism of the lookup in the environment during sub_{var} , sub_l and sub_{if} reductions.
- C' initial segment of C , i.e. $C = C'\langle C'' \rangle$. Symmetric to the first case. \square

PROPOSITION D.11 (OVERHEAD TRANSPARENCY). Let c be a reachable crumble, and let $a \in \{\text{sub}_{\text{var}}, \text{sub}_l, \text{sub}_{\text{if}}\}$. If $c \rightarrow_a d$ then $c_{\downarrow} = d_{\downarrow}$.

PROOF. Let $c := C\langle(b, e_f)\rangle \rightarrow_a C\langle(b', e'_f)\rangle =: d$, and let e'_f, e''_f be such that $e_f = e'_f[x \leftarrow e_f(x)]e''_f$, noting that x does not occur in e''_f by Lemma D.9.1 and Lemma D.9.2. We first prove that $(b, e_f)_{\downarrow} = (b', e'_f)_{\downarrow}$:

- Case sub_{var} , i.e. $b := x$ and $b' = e_f(x)$. By Lemma B.4, $(x, e'_f[x \leftarrow e_f(x)]e''_f)_{\downarrow} = (x, e'_f e''_f)_{\downarrow} \{x \leftarrow (e_f(x), e'_f)_{\downarrow}\} = (e_f(x), e'_f)_{\downarrow}$ as c is well-named (Lemma D.9.1). By Lemma D.9.2, $\text{fv}(e_f(x)) \cap \text{dom}(e'_f[x \leftarrow e_f(x)]) = \emptyset$, therefore $(e_f(x), e'_f)_{\downarrow} = (e_f(x), e_f)_{\downarrow}$, and we conclude with $(x, e_f)_{\downarrow} = (e_f(x), e_f)_{\downarrow}$.
- Case sub_l , i.e. $b := xv$ and $b' = e_f(x)v$. Since $(xv, e_f)_{\downarrow} = (x, e_f)_{\downarrow}(v, e_f)_{\downarrow}$ (Prop. B.8), we can use the point above to conclude.
- Case sub_{if} , i.e. $b := \text{if } x \text{ then } c \text{ else } d$ and $b' = \text{if } e_v(x) \text{ then } c \text{ else } d$. Since $(\text{if } x \text{ then } c \text{ else } d, e_v)_{\downarrow} = \text{if } (x, e_v)_{\downarrow} \text{ then } (c @ e_v)_{\downarrow} \text{ else } (d @ e_v)_{\downarrow}$ (Prop. B.8), we can use the point above to conclude.

We now prove that $C\langle(b, e_f)\rangle_{\downarrow} = C\langle(b', e'_f)\rangle_{\downarrow}$ under the hypothesis that $(b, e_f)_{\downarrow} = (b', e'_f)_{\downarrow}$. By cases on C : if $C := \langle \cdot \rangle$ just use the hypothesis. Otherwise $C :=$

$(b'', e[x \leftarrow \langle \cdot \rangle])$ and so $(b'', e[x \leftarrow b]e_f)_{\downarrow} = (b'', eef)_{\downarrow} \{x \leftarrow (b, e_f)_{\downarrow}\} = (b'', eef)_{\downarrow} \{x \leftarrow (b', e'_f)_{\downarrow}\} = (b'', e[x \leftarrow b']e_f)_{\downarrow}$ by Lemma B.4. \square

LEMMA D.12 (SUBSTITUTION). Let t and u be terms, x_1, \dots, x_n pairwise distinct variables, and f_1, \dots, f_n be fireballs. If $t \rightarrow_a u$ for $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, \text{@e}\}$, then $t\{x_1 \leftarrow f_1, \dots, x_n \leftarrow f_n\} \rightarrow_{c\beta_f} u\{x_1 \leftarrow f_1, \dots, x_n \leftarrow f_n\}$.

PROOF. Let $\sigma := \{x_1 \leftarrow f_1, \dots, x_n \leftarrow f_n\}$. By induction on the definition of $t \rightarrow_{\beta_v} u$. Cases:

- *Top level*.
 - β_v -step, i.e. $t := (\lambda y.s)v' \mapsto_{\beta_v} s\{y \leftarrow v'\} =: u$ and we can suppose without loss of generality that $y \notin \bigcup_{i=1}^n \text{fv}(f_i) \cup \{x_1, \dots, x_n\}$. According to Lemma D.4, $v'\sigma$ is a fireball. As a consequence, $t\sigma = (\lambda y.s\sigma)(v'\sigma) \rightarrow_{\beta_f} s\sigma\{y \leftarrow v'\sigma\} = s\{y \leftarrow v'\sigma\}\sigma = u\sigma$.
 - the *ift, iff, ife, @e* steps are similar to the β_v -step. Lemma D.4 is used in the proof of the *ife*-step to prove that a *ife*-redex where the guard is an abstraction is mapped to a *ife*-redex of the same kind.
- *Application right*, i.e. $t := sr \rightarrow_a sq =: u$ with $r \rightarrow_a q$; by *i.h.* $r\sigma \rightarrow_{c\beta_f} q\sigma$, and therefore $t\sigma = s\sigma(r\sigma) \rightarrow_{c\beta_f} s\sigma(q\sigma) = u\sigma$.
- *Application left*, i.e. $t := sv' \rightarrow_a rv' =: u$ with $s \rightarrow_a r$; by *i.h.*, $s\sigma \rightarrow_{c\beta_f} r\sigma$; according to Lemma D.4, $v'\sigma$ is a fireball and hence $t\sigma = s\sigma(v'\sigma) \rightarrow_{c\beta_f} r\sigma(v'\sigma) = u\sigma$.
- *If-then-else guard*, i.e. $t := \text{if } r \text{ then } s \text{ else } p \rightarrow_a \text{if } q \text{ then } s \text{ else } p =: u$ with $r \rightarrow_a q$; by *i.h.* $r\sigma \rightarrow_{c\beta_f} q\sigma$, and therefore $t\sigma = \text{if } r\sigma \text{ then } s\sigma \text{ else } p\sigma \rightarrow_{c\beta_f} \text{if } q\sigma \text{ then } s\sigma \text{ else } p\sigma = u\sigma$. \square

LEMMA D.13 (SUBSTITUTION OF FIREBALLS). If $e_f := [x_1 \leftarrow b_1] \dots [x_n \leftarrow b_n]$ is a f -environment, then $\sigma(e_f) = \{x_1 \leftarrow u_1, \dots, x_n \leftarrow u_n\}$ where all the u_i 's are fireballs.

PROOF. By induction on the length of e . If $e_f := \epsilon$, then $\sigma(e_f) = \{\}$ and the statement is vacuously true. Otherwise $e_f := [x \leftarrow b]e'_f$ with $e'_f := [x_1 \leftarrow b_1] \dots [x_n \leftarrow b_n]$ and then $\sigma(e_f) = \{x \leftarrow (b, e'_f)_{\downarrow}\} \cup \sigma(e'_f)$; by *i.h.* (since e'_f is a f -environment), $\sigma(e'_f) = \{x_1 \leftarrow u_1, \dots, x_n \leftarrow u_n\}$ where all the u_i 's are fireballs; also $(b, e'_f)_{\downarrow}$ is a fireball by definition of f -environment, as $e_f = E\langle(b, e')\rangle$ with $E := [x \leftarrow \langle \cdot \rangle]$; therefore, $\sigma(e_f) = \{x \leftarrow (b, e'_f)_{\downarrow}, x_1 \leftarrow u_1, \dots, x_n \leftarrow u_n\}$ satisfies the statement. \square

LEMMA D.14 (READ-BACK VS. APPEND). For any crumble c and f -environment e_f , $(c @ e_f)_{\downarrow} = c_{\downarrow} \sigma(e_f)$.

PROOF. By induction on the length of e_f . If $e_f := \epsilon$ then $\sigma(e_f) = \{\}$ and hence $(c @ e_f)_{\downarrow} = c_{\downarrow} = c_{\downarrow} \{\} = c_{\downarrow} \sigma(e_f)$. Otherwise $e_f := [x \leftarrow b]e'_f$ where $x \notin \text{dom}(e'_f)$; by *i.h.* (since e'_f is a fireball environment), $((c @ [x \leftarrow b]) @ e'_f)_{\downarrow} = (c @ [x \leftarrow b])_{\downarrow} \sigma(e'_f)$ and $(b, e'_f)_{\downarrow} = ((b, \epsilon) @ e'_f)_{\downarrow} = (b, \epsilon)_{\downarrow} \sigma(e'_f) = b_{\downarrow} \sigma(e'_f)$; by the definitions of append and read-back, $(c @ [x \leftarrow b])_{\downarrow} = c_{\downarrow} \{x \leftarrow b_{\downarrow}\}$; therefore, $(c @ e_f)_{\downarrow} = ((c @ [x \leftarrow b]) @ e'_f)_{\downarrow} = (c @ [x \leftarrow b])_{\downarrow} \sigma(e'_f) = (c_{\downarrow} \{x \leftarrow b_{\downarrow}\}) \sigma(e'_f) = c_{\downarrow} (\{x \leftarrow b_{\downarrow} \sigma(e'_f)\} \cup \sigma(e'_f)) = c_{\downarrow} \{x \leftarrow (b, e'_f)_{\downarrow}\} \cup \sigma(e'_f) = c_{\downarrow} \sigma(e_f)$. \square

LEMMA D.15. If $c_{\downarrow} \rightarrow_a d_{\downarrow}$ for $a \in \{\beta_v, \text{ift}, \text{iff}, \text{ife}, @e\}$, then $(c @ e_f)_{\downarrow} \rightarrow_{\beta_f} (d @ e_f)_{\downarrow}$.

PROOF. According to Lemma D.13, $\sigma(e_f) = \{x_1 \leftarrow f_1, \dots, x_n \leftarrow f_n\}$ where f_1, \dots, f_n are fireballs. By Lemma D.14 and Lemma D.12, $(c @ e_f)_{\downarrow} = c_{\downarrow} \sigma(e_f) \rightarrow_{\beta_f} d_{\downarrow} \sigma(e_f) = (d @ e_f)_{\downarrow}$. \square

Note that Lemma D.15 does not hold if we let \rightarrow_a be \rightarrow_{β_i} in the hypothesis. Indeed, take $c := ((\lambda x.(x, \epsilon))y, [y \leftarrow (zz, \epsilon)])$ and $d := (y, [y \leftarrow zz])$ and $e_f := [z \leftarrow \lambda x.(xx, \epsilon)]$: then, $c_{\downarrow} = (\lambda x.x)(zz) \rightarrow_{\beta_i} zz = d_{\downarrow}$ but $(c @ e_f)_{\downarrow} = (\lambda x.x)((\lambda x.xx)\lambda x.xx) \not\rightarrow_{\beta_f} (\lambda x.xx)\lambda x.xx = (d @ e_f)_{\downarrow}$. The problem is essentially due to the fact that fireballs, contrary to values, are not closed by substitution: this a notable difference between the closed case (where the normal forms coincide with closed values) and the open case (where the normal forms coincide with fireballs).

PROPOSITION D.16 (PRINCIPAL PROJECTION). Let c be a reachable crumble. If $c \rightarrow_a d$ for $a \in \{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e\}$ then $c_{\downarrow} \rightarrow_a d_{\downarrow}$.

PROOF. Note that for every $b, e_v, (b, e_v) = (b, \epsilon) @ e_v$. Therefore all steps can be written in the form $C\langle(b, e_v)\rangle \rightarrow_a C\langle(c @ e_v)\rangle$ where b is not a crumbled value. The crumble context C unfolds to a right v-context by Lemma D.9.4. We need to prove that $C\langle(b, e_v)\rangle_{\downarrow} \rightarrow_a (C\langle(c @ e_v)\rangle)_{\downarrow}$. By Lemma B.10 and Lemma D.15, it suffices to prove that $C\langle(b, \epsilon)\rangle_{\downarrow} \rightarrow_b C\langle c \rangle_{\downarrow}$ for all $a \neq \beta_f$ and $b = a$ or for $a = \beta_f$ and $b = \beta_v$.

We proceed by cases on the rule \rightarrow_a .

- Rule β_f : we need to prove that $C\langle((\lambda x.c)v, \epsilon)\rangle_{\downarrow} \rightarrow_{\beta_v} C\langle e @ [y \leftarrow v] \rangle_{\downarrow}$ where $e @ [y \leftarrow v] := (c @ [x \leftarrow v])^\alpha$.

$$\begin{aligned} C\langle((\lambda x.c)v, \epsilon)\rangle_{\downarrow} &= C_{\downarrow}\langle(\lambda x.c_{\downarrow})v_{\downarrow}\rangle && \text{by Corollary B.11.1} \\ &=_{\alpha} C_{\downarrow}\langle(\lambda y.e_{\downarrow})v_{\downarrow}\rangle \\ &\rightarrow_{\beta_v} C_{\downarrow}\langle e_{\downarrow}\{y \leftarrow v_{\downarrow}\}\rangle \\ &= C_{\downarrow}\langle e @ [y \leftarrow v] \rangle_{\downarrow} \\ &= C\langle e @ [y \leftarrow v] \rangle_{\downarrow} && \text{by Corollary B.11.1} \\ &= C\langle(c @ [x \leftarrow v])^\alpha \rangle_{\downarrow}. \end{aligned}$$

Note that the second use of Corollary B.11.1 requires that $C \# (e @ [y \leftarrow v])$ i.e. that $\text{fv}(C) \cap \text{dom}(e @ [y \leftarrow v]) = \emptyset$, which follows from the side condition about α -renaming in the β_f rule.

- Rules $\text{ift}, \text{iff}, \text{ife}, @e$: a quick check by cases over $(b, e_v) \mapsto_a c @ e_v$ shows that $b_{\downarrow} \rightarrow_a c_{\downarrow}$.

For example, $(\text{if true then } c \text{ else } d, e_v) \mapsto_{\text{ift}} c @ e_v$ and $(\text{if true then } c \text{ else } d)_{\downarrow} = \text{if true then } c_{\downarrow} \text{ else } d_{\downarrow} \rightarrow_{\text{ift}} c_{\downarrow}$. The other cases are all similar.

Thus

$$\begin{aligned} C\langle(b, \epsilon)\rangle_{\downarrow} &= C_{\downarrow}\langle b_{\downarrow} \rangle && \text{by Corollary B.11.1} \\ &\rightarrow_a C_{\downarrow}\langle c_{\downarrow} \rangle \\ &= C\langle c \rangle_{\downarrow} && \text{by Corollary B.11.1} \end{aligned}$$

Note that the second use of Corollary B.11.1 requires that $C \# c$. The property holds because all substitutions outside abstractions in the rhs of the rules \rightarrow_a under consideration were such in the lhs, and because reachable crumbles are well-named (Lemma D.9.1). \square

LEMMA D.17 (HALT). If c is oCr-normal then c_{\downarrow} is $\rightarrow_{c\beta_f}$ -normal.

PROOF. By Prop. 6.3, if c is normal then it is a f -crumble i.e. $c = c_f$. By definition of c_f , c_{\downarrow} is a fireball. By harmony for $\lambda_{\text{fire}}^{\text{if}}$, c_{\downarrow} is $c\beta_f$ -normal. \square

PROOF OF (THM. 6.5). Let c be a crumble that is reachable by the Open Crumble GLAM.

- (1) Initialization: $t_{\downarrow} = t$
- (2) Principal projection: if $c \rightarrow_a d$ then $c_{\downarrow} \rightarrow_a d_{\downarrow}$ for $a \in \{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e\}$.
- (3) Overhead transparency: if $c \rightarrow_a d$ then $c_{\downarrow} = d_{\downarrow}$ for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}\}$.
- (4) Determinism: the transition \rightarrow_{oCr} is deterministic.
- (5) Halt: if c is \rightarrow_{oCr} -normal then c_{\downarrow} is $\rightarrow_{c\beta_f}$ -normal.
- (6) Overhead termination: \rightarrow_a terminates, for any rule $a \in \{\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}\}$.

Therefore, the Open Crumble GLAM, the right-to-left conditional fireball evaluation $\rightarrow_{c\beta_f}$ and the read-back $(\cdot)_{\downarrow}$ form an implementation system.

PROOF. (1) See Prop. 4.2.

(2) See Prop. D.16.

(3) See Prop. D.11.

(4) See Lemma D.10.

(5) See Lemma D.17.

(6) Immediate consequence of the equivalent of Lemma 5.8 for the Open Crumble GLAM. \square

D.3 The Open Pointed Crumble GLAM

The design of a machine, called *Open Pointed Crumble GLAM*, for the open case that decomposes the search for the next redex and (un)plugging in $O(1)$ transitions follows the same pattern as for the Pointed Crumble GLAM in Appendix C.4 (p. 22) for the closed case. In particular, the definitions of pointed environment, translation $\iota(\cdot)$ (from crumbles to pointed environments) and read-back $(\cdot)_{\downarrow}$ (from pointed environments to crumbles) are the same. The transitions of the Open Pointed Crumble GLAM (whose union is noted $\rightarrow_{\text{poCr}}$) differ from the ones of the Pointed Crumble GLAM exactly as the transitions of Open Crumble GLAM differ from the ones of the Crumble GLAM.

D.3.1 Implementation.

Pointed environments. For the Open Pointed Crumble GLAM, the definitions of pointed environment, translation $\iota(\cdot)$ (from crumbles to pointed environments) and read-back $(\cdot)_{\downarrow}$ (from pointed environments to crumbles) are the same as for the Pointed Crumble GLAM, see Appendix C.4 (p. 22).

Evaluation. The transition rules of the Open Pointed Crumble GLAM are:

$$\begin{aligned}
& e[x \leftarrow (\lambda y. c) v] \mid e_f \rightarrow_{\beta_f} e[x \leftarrow b]e'[z \leftarrow v] \mid e_f \\
& e[x \leftarrow \text{if true then } c \text{ else } d] \mid e_f \rightarrow_{\text{ift}} e[x \leftarrow b]e' \mid e_f \\
& e[x \leftarrow \text{if false then } c \text{ else } d] \mid e_f \rightarrow_{\text{iff}} e[x \leftarrow b]e' \mid e_f \\
& e[x \leftarrow \text{if } v \text{ then } c \text{ else } d] \mid e_f \rightarrow_{\text{ife}} e[x \leftarrow \text{err}] \mid e_f \\
& e[x \leftarrow v w] \mid e_f \rightarrow_{@e} e[x \leftarrow \text{err}] \mid e_f \\
& e[y \leftarrow x] \mid e_f \rightarrow_{\text{sub}_{var}} e[y \leftarrow e_f(x)] \mid e_f \\
& e[y \leftarrow x v] \mid e_f \rightarrow_{\text{sub}_l} e[y \leftarrow e_f(x) v] \mid e_f \\
& e[y \leftarrow \text{if } x \text{ then } c \text{ else } d] \mid e_f \rightarrow_{\text{sub}_{if}} e[y \leftarrow \text{if } e_f(x) \text{ then } c \text{ else } d] \mid e_f \\
& e[x \leftarrow b] \mid e_f \rightarrow_{\text{sea}} e \mid [x \leftarrow b]e_f
\end{aligned}$$

where

- i. $\lambda z.(b, e') := (\lambda y. c)^\alpha$ such that $(e[x \leftarrow b]e'[z \leftarrow v] \mid e_f)$ is well-named.
- ii. where $c := (b, e')$
- iii. where $d := (b, e')$
- iv. if $v = \lambda x. e$ or $v = \text{err}$.
- v. if $v \in \{\text{true}, \text{false}, \text{err}\}$.
- vi. if $x \in \text{dom}(e_f)$.
- vii. if none of the other rules is applicable, i.e. when b is a practical value $v \rightarrow x$ or when b is x, xv , or if x then c else d but x is not defined in e_f .

The transition function $\rightarrow_{\text{poCr}}$ of the Open Pointed Crumble GLAM is then defined as the union of the rules above. A *principal transition* of the Open Pointed Crumble GLAM is a transition \rightarrow_a for any rule $a \in \{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e\}$.

Definition D.18 (Reachable state). A pointed environment is said to be *reachable* (in the Open Pointed Crumble GLAM) if it is obtained from evaluation steps starting from the translation $\iota(c)$ of a well-named crumble c .

LEMMA D.19 (INVARIANTS). *Let e_1 be a reachable pointed environment:*

- (1) Freshness: e_1 is well-named;
- (2) Rightmost: $e_1 = (e \mid e_f)$ for some e and some f -environment e_f .

PROOF. By induction on the length of the evaluation sequence leading to e_1 . The base cases hold by the definition of reachability and by the definition of $\iota(\cdot)$. As for the inductive step, we proceed by cases on the transitions:

- (1) For β_f the claim follows from the side condition. The rules in $\{\text{ift}, \text{iff}, \text{ife}, @e, \text{sea}\}$ do not increase the number of explicit substitutions occurring in e_1 outside of abstractions, hence the claim follows from the *i.h.* The rules in $\{\text{sub}_{var}, \text{sub}_l, \text{sub}_{if}\}$ copy a practical value from the environment e_f : note that $e_f(x)$ is either an abstraction (which does not influence well-namedness) or a boolean or an error (which do not contain explicit substitutions).
- (2) The rules in $\{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e, \text{sub}_{var}, \text{sub}_l, \text{sub}_{if}\}$ do not change the evaluated part, hence the claim follows from the

i.h. As for the sea rule, by Lemma D.21 it suffices to prove that (b, e_f) is a f -crumble, knowing that the other transition rules cannot be applied: this follows from Lemma D.7. \square

- (ii) LEMMA D.20 (HARMONY FOR THE OPEN POINTED CRUMBLE GLAM). *Let e_1 be a reachable pointed environment in the Open*
- (iii) *Pointed Crumble GLAM: e_1 is normal if and only if it has the form*
- (iv) *$(\epsilon \mid e_f)$ for some non-empty f -environment e_f .*
- (v) PROOF. The proof of the implication from right to left is trivial. As for the other direction, let e_1 be a reachable normal
- (vi) *pointed environment. By Lemma D.19, e_1 has the form $(e \mid e_f)$.*
- (vi) *Note that e must be empty; otherwise one of the transitions in*
- (vi) *$\{\beta_f, \text{ift}, \text{iff}, \text{ife}, @e, \text{sub}_{var}, \text{sub}_l, \text{sub}_{if}, \text{sea}\}$ could be applied, contradicting the hypothesis that e_1 is normal. \square*
- (vii)

LEMMA D.21. *If $[x \leftarrow b]e$ is a f -environment, then (b, e) is a f -crumble.*

PROOF. Assume that $[x \leftarrow b]e$ is a f -environment, and let $(b, e) = C\langle(b', e')\rangle$ for some C, b', e' . Then $[x \leftarrow b]e = E\langle C\langle b', e' \rangle \rangle$ for $E := [x \leftarrow \langle \cdot \rangle]$. The requirements for a f -crumble follow from the definition of f -environment for $[x \leftarrow b]e$. \square

THEOREM D.22 (IMPLEMENTATION). *Let e_1 a pointed environment that is reachable by the Open Pointed Crumble GLAM.*

- (1) Initialization: $(\iota(c))_{\parallel} = c$ for every crumble c .
- (2) Principal Projection: if $e_1 \rightarrow_a e'_1$ then $(e_1)_{\parallel} \rightarrow_a (e'_1)_{\parallel}$ for any rule $a \neq \text{sea}$.
- (3) Overhead Transparency: if $e_1 \rightarrow_{\text{sea}} e'_1$, then $(e_1)_{\parallel} = (e'_1)_{\parallel}$.
- (4) Determinism: the transition function $\rightarrow_{\text{poCr}}$ is deterministic.
- (5) Halt: if e_1 is normal, then $(e_1)_{\parallel}$ is normal.
- (6) Overhead Termination: \rightarrow_{sea} terminates.

Therefore, the Open Pointed Crumble GLAM (with its transition function $\rightarrow_{\text{poCr}}$), the Open Crumble GLAM (with \rightarrow_{oCr}), and the read-back $(\cdot)_{\parallel}$ form an implementation system.

PROOF.

- (1) Let $c = (b, e)$: by the definitions, $(\iota(c))_{\parallel} = ([x \leftarrow b]e \mid \epsilon)_{\parallel} = (b, e\epsilon) = c$.
- (2) There is a clear one-to-one correspondence between the transitions of the Open Pointed Crumble GLAM and the Open Crumble GLAM (apart from sea). The proof is similar to the one of Thm. C.15.
- (3) By inspection of the rule sea. We need to prove that $(e[x \leftarrow b] \mid e_f)_{\parallel} = (e \mid [x \leftarrow b]e_f)_{\parallel}$. By cases on the structure of e : if $e = \epsilon$, then $(\epsilon[x \leftarrow b] \mid e_f)_{\parallel} = (b, e_f) = (\epsilon \mid [x \leftarrow b]e_f)_{\parallel}$. If instead $e = [y \leftarrow b']e'$, then $(e[x \leftarrow b] \mid e_f)_{\parallel} = (b', e'[x \leftarrow b]e_f) = (e \mid [x \leftarrow b]e_f)_{\parallel}$.
- (4) The rule sea can be applied by definition only when the other rules cannot be applied. The remaining rules apply to a pointed environment of the form $(e[x \leftarrow b] \mid e_f)$, for distinct shapes of b :
 - The rules $\{\beta_f, @e, \text{sub}_l\}$ apply when b is an application, and respectively the application of an abstraction to a crumbled value, of a boolean/error to a crumbled value, and of a variable to a crumbled value.
 - The rule sub_{var} applies when b is a variable.

- The rules in $\{\text{ift}, \text{iff}, \text{ife}, \text{sub}_{\text{if}}\}$ apply when b is a if-then-else, and in clearly disjoint cases according to the structure of the condition: respectively, when the condition is true, false, an abstraction or an error, and a variable.
 - Finally, note that no single rule in $\{\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}\}$ can transition to different pointed environments due to the lookup of a variable which has multiple occurrences in the environment: the lookup is deterministic because the environment is well-named (Lemma D.19.1).
- (5) By Lemma D.20, e_I is normal iff it has the form $(\epsilon \mid e_f)$ for some non-empty f -environment e_f . Then $(e_I)_{\Downarrow} = (\epsilon \mid e_f)_{\Downarrow}$ which is a f -crumble by Lemma D.21. By Prop. 6.3, $(e_I)_{\Downarrow}$ is normal.
- (6) Immediate consequence of forthcoming Cor. D.24 (proved independently). \square

D.3.2 Complexity. The complexity analysis for the Open Pointed Crumble GLAM is analogue to the one for the Pointed Crumble GLAM. More precisely, following the same approach and notations as in Appendix C.4 and in Appendix C.5, we can prove:

LEMMA D.23 (NUMBER OF sea-TRANSITIONS). *Let c be a well-named crumble, and let $\rho: \iota(c) \rightarrow_{\text{poCr}}^* e_I = (\epsilon \mid e_v)$ an execution in the Open Pointed Crumble GLAM. Then $|e|_{\text{len}} \leq |c|_{\text{len}} + (|\rho|_{\beta_f} + |\rho|_{\text{ift}} + |\rho|_{\text{iff}}) \cdot L(c) - |\rho|_{\text{sea}}$.*

COROLLARY D.24. *Let t be a term. For a normalizing execution ρ in the Open Pointed Crumble GLAM starting from $\iota(t)$, we have $|\rho|_{\text{sea}} \leq (|\rho|_{\text{p}} + 1) \cdot |t|$.*

THEOREM D.25 (THE OPEN POINTED CRUMBLE GLAM IS BILINEAR). *For any term t and any Open Pointed Crumble GLAM execution $\rho: \iota(t) \rightarrow_{\text{poCr}}^* e_I$, the cost of implementing ρ on a RAM is $O((|\rho|_{\text{p}} + 1) \cdot |t|)$.*

PROOF. Essentially the same proof as the one of Thm. C.18. \square

Analogously to the closed case, the previous theorem also has a consequence for the non-pointed case. Since the theorem shows that the cost of searching for redexes and the (un)plugging operations can indeed be realised in bilinear time, it is also true that the Open Crumble GLAM can be implemented in bilinear time, improving Thm. 6.6 by removing the *up to search and (un)plugging* side condition.

COROLLARY D.26 (THE OPEN CRUMBLE GLAM IS BILINEAR). *For any term t and any Open Crumble GLAM execution $\rho: \underline{t} \rightarrow_{\text{oCr}}^* c$, the cost of implementing ρ on a RAM is $O((|\rho|_{\text{p}} + 1) \cdot |t|)$.*

E IMPLEMENTATION IN OCAML

The goal of this section is implementing in OCaml the Pointed Crumble GLAM and Open Pointed Crumble GLAM presented in Appendix C.4 (p. 22) and Appendix D.3 (p. 29). We are going to describe the abstract requirements of the data structures, before picking a concrete implementation. The two machines share the same data structures and auxiliary functions, and only differ by five lines of code.

Data structures. The machine works on pointed environments of the form $(e \mid e_v)$. The unevaluated part of the environment e and the evaluated part e_v are subject to different requirements:

- The unevaluated part e is extended only on the right by the transitions $\beta_v, \text{ift}, \text{iff}$ by concatenating an unevaluated environment. Only the rightmost entry is inspected by every reduction rule. The sea transition removes the rightmost entry, moving it to the evaluated part. Therefore the unevaluated environment must implement the *catenable stack* interface, allowing to perform catenation, pop, and topmost inspection in constant time.
- As for the evaluated part e_v , transitions never exploit the sequential structure of e_v . On the contrary, the $\text{sub}_{\text{var}}, \text{sub}_I, \text{sub}_{\text{if}}$ transitions need to access the entry associated with a variable x in constant time. The only other operation required (by the sea rule) is to add an entry to it in constant time.

To satisfy lookup in constant time for e_v , we implement e_v as a store, thus ignoring its list structure. In turn, this choice impacts on the data structure for terms, because it forces (occurrences of) variables to be implemented as pointers to memory locations. More explicitly, an entry $[x \leftarrow b]$ is represented as a node n which is a record containing a field content holding b together with additional fields soon to be described. If α is the address of n , occurrences of x are represented in OCaml as $\text{Shared}(n)$, which means a memory cell tagged as Shared and holding the pointer α .

Occurrences of λ -bound variables are instead presented as $\text{Var}(v)$ where v is a unique identifier for that variable. Thus the data structure for terms is:

```
...
and term =
  | Err
  | True
  | False
  | Var of var
  | Lam of var * crumblep
  | App of term * term
  | IfThenElse of term * crumblep * crumblep
  | Shared of node
...
```

A crumble is a term coupled with an unevaluated environment. As discussed above, the unevaluated environment implements a catenable stack interface. The simplest implementation, which we adopted, is a linked list of nodes, referenced by two pointers to the first and the last entries of the stack, or a special value to represent the empty stack. Concretely, in OCaml we use an option type for that.

```
...
and env = (node * node) option
and crumblep = term * env
...
```

Each node has a field *prev* used to point to the previous entry in the environment. An additional field *copying* and *mutability* of

```

type var = { dummy : unit }
type node =
  { mutable content : term
    ; mutable copying : bool
    ; mutable prev : node option }
and env = (node * node) option
and crumblep = term * env
and term =
  | Err
  | True
  | False
  | Var of var
  | Lam of var * crumblep
  | App of term * term
  | IfThenElse of term * crumblep * crumblep
  | Shared of node

let mk_node content =
  { content ; copying = false ; prev = None }

let push n e =
  n.prev <- Some e

```

Table 1: Data structures

all fields are required to implement α -renaming in linear time; we explain their use later. Therefore a node is:

```

type node =
  { mutable content : term
    ; mutable copying : bool
    ; mutable prev : node option }
...

```

Unreachable nodes can be garbage-collected by the runtime of OCaml. Because the evaluator holds a pointer to the unevaluated environment, only evaluated nodes can be garbage-collected.

Finally, we implement the datatype of unique λ -bound variable identifiers `var` as the address of an OCaml record that holds no useful information. Thus comparing variables can be achieved using pointer equality `==`. Concrete implementations can add fields to the record, for example to associate the name of the variable as a string.

```
type var = {dummy: unit} (*no empty records in OCaml*)
```

A summary of the data structures can be found in Table 1.

Implementation of transitions. The code that implements evaluation in the closed and open cases can be found in Table 4. For the same reason as it is discussed in Appendix C.4 (p. 22), the input of the evaluation functions `eval_c/eval_o` is not a crumble, but a crumbled unevaluated environment $\iota(b, e) = [x \leftarrow b]e$ where x is a fresh variable. More precisely, the evaluation functions take in input just a node n , which is the rightmost entry of the crumbled unevaluated environment that has to be evaluated.

The code that implements the transition β_v for $(\lambda y.c)v$ is the most complex because it must:

- (1) α -rename $\lambda y.c$ to $\lambda y'.(b', e')$;
- (2) change the top of the unevaluated environment (stack) from $[n \leftarrow (\lambda y.c)v]$ to $[n \leftarrow b']$ and append e' to it;
- (3) push $[y' \leftarrow v]$ on top of the unevaluated environment.

To implement the previous steps efficiently, the code creates the node y' containing v and then calls a function `copy_crumbp y y' c` that performs step 1 in linear time, returning the new unevaluated crumble c' . Then $n @ c'$ performs step 2 in constant time by concatenating c' to the unevaluated environment (and therefore to n , its topmost element). The code of `__ @` is given in Table 2.

Finally, push $y' (n @ c')$ performs step 3 in constant time by pushing y' on the new top of the unevaluated environment.

The transitions `subI`, `subvar`, `subif`, `@e`, `ife` rules just update the content of the top of the unevaluated environment in the required way.

The transitions `ift` and `iff` perform in constant time the plugging of the crumble c to the unevaluated environment, then call the evaluation function on the new topmost entry of the unevaluated environment.

The transition `sea` is implemented by the function `pop` that pops the top of the unevaluated environment and calls evaluation on the new top (if present):

```

...
and pop n =
  match n.prev with
  | None -> n.content
  | Some p ->
    n.prev <- None;
    eval_c/o p
...

```

When a node is popped, its `prev` pointer is unset to facilitate garbage-collection of unreferenced nodes.

Otherwise a normal form is reached, and evaluation returns the term that, pointing to the evaluated environment, consists of the normal crumble.

Let us remark that the implementation is tail-recursive⁴; since OCaml optimizes tail-recursion, the machine only consumes constant space on the process execution stack.

As a minor optimization to the expected code, our implementation merges execution of rule `subvar` with that of the `sea` step which always follows the former. The merging is obtained calling `pop` in place of `eval_c/eval_o`.

The complexity of each case is $O(1)$, but for the β_v rule which requires a renaming of crumble (`copy_crumbp`). It remains to see how this operation can be implemented with linear complexity.

Implementation of α -renaming. We implement α -renaming of unevaluated environments by creating a copy of the environment. The representation in memory of the environment is a DAG because terms in the nodes of the environment contain occurrences of Shared nodes defined in the same environment. Therefore we need to implement a copy algorithm over DAGs that runs in linear time.

⁴when the `pop` function is inlined

<pre> let rec eval_c n = match n.content with App(Lam(y,c), t) -> (* β_v *) let y' = mk_node t in let c' = copy_crumbp y y' c in push y' (n @ c') ; eval_c y' App (Shared{content=t1}, t2) -> (* sub_I *) n.content <- App(t1, t2); eval_c n App((True False Err), _) -> (* @e *) n.content <- Err ; eval_c n IfThenElse(True,c,_) (* ift *) IfThenElse(False,_,c) (* iff *) -> eval_c (n @ c) IfThenElse (Shared {content=t1}, t2,t3) -> (* sub_{if} *) n.content <- IfThenElse(t1,t2,t3) ; eval_c n IfThenElse((Lam _ Err),_,_) -> (* ife *) n.content <- Err ; eval_c n Shared {content} -> (* sub_{var} *) n.content <- content ; pop n Lam _ Err True False -> (* sea *) pop n Var _ App(Var _, _) IfThenElse(Var _,_,_) -> failwith "Open_term" _ -> assert false </pre>	<pre> let rec eval_o n = match n.content with App(Lam(y,c), t) -> (* β_f *) let y' = mk_node t in let c' = copy_crumbp y y' c in push y' (n @ c') ; eval_o y' App (Shared{content= (Lam _ True False Err) as t1}, t2) -> (* sub_I *) n.content <- App(t1, t2); eval_o n App((True False Err), _) -> (* @e *) n.content <- Err ; eval_o n IfThenElse(True,c,_) (* ift *) IfThenElse(False,_,c) (* iff *) -> eval_o (n @ c) IfThenElse (Shared{content= (Lam _ True False Err) as t1}, t2,t3) -> (* sub_{if} *) n.content <- IfThenElse(t1,t2,t3) ; eval_o n IfThenElse((Lam _ Err),_,_) -> (* ife *) n.content <- Err ; eval_o n Shared{content= (Lam _ True False Err) as c} -> (* sub_{var} *) n.content <- c ; pop n Lam _ Err True False -> (* sea *) pop n Var _ App(Var _, _) Shared _ App(Shared _, _) IfThenElse (Var _,_,_) -> (* sea *) pop n _ -> assert false </pre>
--	---

Figure 4: Evaluation: closed (left) vs open (right)

The algorithm consists in using the content field of nodes to perform the renaming. When a node is being copied, it is temporarily put in the copying=true status, and its content field is changed to point to the corresponding new node. Then, the rest of the environment is copied recursively. When an occurrence of a node that

is being copied is found in the term being copied, it is replaced with the new node stored in the content field of the old one. Finally, when the copy is over, the copying status of every node is reset to false and the previous value of content is restored, yielding the original environment.

The auxiliary copying_node $y \ y' \ f$ function, where y is the node to be copied to y' , implements the idea above by temporarily putting y in copying=true status, until f is executed.

```
let copying_node y y' f =
  let saved = y.content in
  y.content <- Shared y' ;
  y.copying <- true ;
  let res = f () in
  y.content <- saved ;
  y.copying <- false ;
  res
```

The `copy_crumbp v n p` function (Table 2) not only implements the algorithm above by copying p , but it also replaces occurrences of `Var v` (the bound variable in $a \rightarrow \beta_v$ redex) with `Shared n` (the new node pointing to the argument of the redex), i.e.: let $c := b_0, [x_1 \leftarrow b_1] \dots [x_k \leftarrow b_k]$; then `copy_crumbp v n p` = $b'_0, [y'_1 \leftarrow b'_1] \dots [y'_k \leftarrow b'_k]$, where $b'_i := b_i \{v \leftarrow n\} \{x_{i+1} \leftarrow x'_{i+1}\} \dots \{x_k \leftarrow x'_k\}$.

The implementation of `copy_crumbp v n p` uses an auxiliary function `copy_env` that iterates over the environment copying it. The function `copy` copies a bite.

Crumbling. The code in Table 3 takes an unevaluated environment and returns the corresponding crumbled unevaluated environment. It generalizes the function `_` in the paper that turns terms into crumbles. It mainly consists of three mutually recursive functions:

- (1) `aux_term c e` translates c to $d = (b', e')$, appends e' to e and returns the obtained crumble.
- (2) `aux_value c e` checks whether c is a value, computing either `̄`, or `_` by calling `aux_term c e`.
- (3) `aux_env c e` creates a copy of (c, e) in crumbled form in linear time reusing the same trick of the copying flag as in α -renaming.

```
let copy_crumbp v n p =
  let rec copy = function
    | Var v' when v == v' -> Shared n
    | Shared {content; copying} when copying -> content
    | Err | True | False | Var _ | Shared _ as c -> c
    | App(c1,c2) -> App(copy c1, copy c2)
    | IfThenElse(c,p,q) ->
      IfThenElse(copy c, copy_crumbp p, copy_crumbp q)
    | Lam(v,e) -> Lam(v, copy_crumbp e)
  and copy_env c e =
    let n' = mk_node (copy e.content) in
    copying_node e n' (fun () ->
      match e.prev with
      | None -> copy c, n'
      | Some prev ->
        let c', b', e' = copy_env c prev in
        push n' e' ;
        c', b', n')
    and copy_crumbp (c,e) =
      match e with
      | None -> copy c, None
      | Some (b,e) ->
        let c', b', e' = copy_env c e in
        c', Some (b', e')
    in copy_crumbp p

let (@) n (c, env) =
  n.content <- c ;
  match env with
  | None -> n
  | Some (b,e) -> push b n ; e
```

Table 2: The `copy_crumbp` function

```

let dummy = Var (mk_var ())

let iota e =
  let star = mk_node dummy in
  star @ e

let anf p =
  let rec aux_term c e = match c with
  | Var _ | Err | True | False -> c, e
  | App(v, w) ->
    let v, e = aux_val v e in
    let w, e = aux_val w e in
    App(v, w), e
  | IfThenElse(v, p, q) ->
    let c, e = aux_val v e in
    let p = aux_crumbp p in
    let q = aux_crumbp q in
    IfThenElse(c,p,q), e
  | Lam(x, p) -> Lam(x, aux_crumbp p), e
  | Shared n -> n.content, e
  and aux_val c e = match c with
  | App _ | IfThenElse _ ->
    let n = mk_node dummy in
    let b =
      (match e with
        None -> n
        | Some (b,e) -> push n e ; b) in
    let c, e = aux_term c (Some (b,n)) in
    n.content <- c;
    Shared n, e
  | Var _ | Lam _ | Shared _
  | Err | True | False -> aux_term c e
  and aux_env c e =
    let p = aux_term e.content None in
    match e.prev with
    | None -> aux_term c (snd p)
    | Some prev ->
      let n = mk_node dummy in
      let last = n @ p in
      let (c,env) =
        copying_node e n (fun () -> aux_env c prev) in
      (match env with
        None -> c, Some(n,last)
        | Some (b,e) -> push n e ; c, Some (b,last))
    and aux_crumbp (c,env) =
      match env with
      None -> aux_term c None
      | Some (_,e) -> aux_env c e in
    iota (aux_crumbp p)

```

Table 3: The anf function