Recurrence Extraction for Functional Programs through Call-by-Push-Value

G. A. KAVVOS*, EDWARD MOREHOUSE, DANIEL R. LICATA, and NORMAN DANNER, Wesleyan University, United States of America

The main way of analyzing the complexity of a program is that of extracting and solving a recurrence that expresses its running time in terms of the size of its input. We develop a method that automatically extracts such recurrences from the syntax of higher-order recursive functional programs. The resulting recurrences, which are programs in a call-by-name language with recursion, explicitly compute the running time in terms of the size of the input. In order to achieve this in a uniform way that covers both call-by-name and call-by-value evaluation strategies, we use Call-by-Push-Value (CBPV) as an intermediate language. Finally, we use domain theory to develop a denotational cost semantics for the resulting recurrences.

CCS Concepts: • **Theory of computation** → **Program verification**; **Program analysis**; *Denotational semantics*; • **Software and its engineering** → *Functional languages*.

Additional Key Words and Phrases: recurrence extraction, resource analysis, cost semantics, higher order recurrences, denotational semantics, call by push value, general recursion

ACM Reference Format:

G. A. Kavvos, Edward Morehouse, Daniel R. Licata, and Norman Danner. 2020. Recurrence Extraction for Functional Programs through Call-by-Push-Value. *Proc. ACM Program. Lang.* 4, POPL, Article 15 (January 2020), 31 pages. <https://doi.org/10.1145/3371083>

1 INTRODUCTION

Functional programmers typically analyze time, space, or other resource usage of their programs using the *extract-and-solve* method. First, we *extract a recurrence* from the program. In this context, a recurrence is a mathematical object—usually an inequality—that expresses an upper bound for the running time of a program in terms of the *size* of its input. Depending on the task at hand, this notion of size may vary. For example, if the input is a tree, we may define its size to be its number of nodes, its depth, or some more complicated expression. The second step consists of *solving* this recurrence: mathematical methods are used to express it (or a suitably looser version of it) in a non-recursive *closed form* and big-*O* bound. While this method is taught in introductory textbooks (e.g. Bird [2014, §7]), there is no *formal* connection between the program and the extracted recurrence. In this work we concentrate on the first of those steps: we seek a method to *automatically extract a recurrence* from the syntax of a recursive functional program, in such a way that we can prove a formal *bounding theorem* relating the extracted recurrence to the program’s operational cost.

*Current affiliation: Department of Computer Science, Aarhus University

Authors’ address: G. A. Kavvos, g.a.kavvos@gmail.com; Edward Morehouse, emorehouse@wesleyan.edu; Daniel R. Licata, dlicata@wesleyan.edu; Norman Danner, ndanner@wesleyan.edu, Department of Mathematics and Computer Science, Wesleyan University, 265 Church Street, Middletown, CT, 06459, United States of America.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/1-ART15

<https://doi.org/10.1145/3371083>

Fig. 1. Recurrence extraction

Danner et al. [2015] present such a method for a call-by-value terminating language with inductive types and structural recursion. Their method consists of the following steps:

- Given a program M in the *source language*—i.e. the language that we wish to analyse—a *syntactic recurrence* $\llbracket M \rrbracket$ is extracted. This recurrence is expressed in an appropriate *recurrence language*, which includes primitives for expressing cost. This step is close in spirit to a monadic translation of the program into the writer monad. This approach is able to accommodate higher-order functions, assigning to them a higher-order recurrence expressing their cost in terms of a recurrence for their input.
- A *bounding relation* between source programs and syntactic recurrences is defined by induction on the types of the source language, i.e. as a logical relation. Intuitively, a source program is bounded by a recurrence if the components of the recurrence express *upper bounds* for the attributes of the source program, e.g. evaluation cost, size, etc.
- Following that, a *bounding theorem* is proved. This shows that every source program M is bounded by the extracted recurrence $\llbracket M \rrbracket$.
- Finally, a denotational semantics is provided for the recurrence language. Depending on the intended application, this semantics abstracts inductive data types to some notion of size. For example, to consider binary trees up to their height we might interpret them as natural numbers, with the node constructor interpreted as the maximum. Alternatively, the node constructor may be interpreted as addition, thus yielding the number of nodes.

This strategy is shown schematically in Figure 1.

In this previous work, the *recurrence language* was taken to be a call-by-name language, so as to be as ‘close to mathematics’ as possible. Composing the extraction with the semantic interpretation then yields what one might call a *semantic recurrence*, which is intended to match the recurrence we would informally write when teaching undergraduate students—at least in the context of first-order programs. Thus the entire process gives a formal account and justification for informal cost analysis techniques. Factoring this into a syntactic and a semantic step is a useful tool for obtaining a simplifying account, which is additionally modular in the different notions of size for inductive data types. Nevertheless, it is still somewhat rigid with respect to changes to the *source language*: each source language requires a new and complex logical relations proof.

In this paper, we improve upon the above approach in several ways:

- (1) We factor the syntactic phase into a cost-preserving embedding of the source language in an *intermediate language*, followed by a recurrence extraction for the intermediate language in the style of Danner et al. [2015]. The bounding theorem is proved once for the intermediate language, so for each source language we need only prove that the cost-preserving embeddings are correct. This strategy is illustrated Figure 2.
- (2) In order to support a wide variety of languages, we choose as an intermediate language Call-by-Push-Value (CBPV) [Levy 2003], which embeds both call-by-value and call-by-name evaluation. We define recurrence extraction and prove a bounding theorem for CBPV, and obtain recurrence extraction and bounding theorems for both call-by-value (CBV) and call-by-name (CBN) source languages by embedding them in CBPV.

Fig. 2. Recurrence extraction through CBPV

- (3) Our intermediate language is a version of CBPV with general recursion. Thus, we obtain a new recurrence extraction and bounding theorem result for CBV with general recursion, and a simpler recurrence extraction and bounding theorem for CBN than Kim [2016]. General recursion is important in our setting for writing divide-and-conquer algorithms in the standard way: in order to give a formal analogue of the standard informal techniques, we should not require the programmer to embed a termination metric in the source program.
- (4) We generalize the denotational semantics from sets equipped with a size ordering to domains equipped with a size ordering. This involves some subtle questions regarding the interaction of the size ordering with the information order of the domains. We develop a semantics that identifies the *bottom* element of the domains (non-termination) with the *maximum* element of the size ordering (infinity), and more generally reverse-includes the information order in the size order—a more defined recurrence is a more precise bound. The key features of this semantics are encoded in a syntactic recurrence language PCFc.

In more detail, our technical development proceeds as follows. We first define three variants of PCF, starting in Section 2 with relatively standard CBV and CBN source languages. To isolate the issues involved in considering general recursion, we choose source languages with only natural numbers as data (treated as a flat base type). Next, in Section 3 we define the recurrence language PCFc, a version of call-by-name PCF with an additional type of costs. The key axiomatic component of PCFc is its *size order* (really, a pre-order relation). The size order codifies the minimum requirements necessary to prove the bounding theorem, while subsequent denotational interpretations of PCFc can refine this to specific notions of cost and size. The primary conceptual difficulty in the size order is understanding how it interacts with the implicit information order that arises from models of PCF, which is necessary for verifying the bounding theorem for recurrences extracted from recursively-defined functions. Our approach takes more defined recurrences to be *smaller* in the size order, i.e. to be better bounds. This is encoded in two of the rules that axiomatize the size order: one that asserts that rational chains are decreasing in the size order, and one that asserts that if all the approximants to a fixed point are a bound, then so is the fixed point itself.

Next, we describe the “end-user” recurrence extraction functions and corresponding bounding theorems for CBV and CBN source languages in Section 4. The statements of these theorems for the two source languages are independent of the intermediate language, but we will prove them by

factoring them through it, as indicated in Figure 2. These two examples illustrate the flexibility of our approach: even though they originate from the same general theorem, the two bounding relations are very different in spirit.

The rest of the paper is largely devoted to proving that the extraction functions in Section 4 are correct using the approach of Figure 2. We define the version of Levy’s *Call-by-Push-Value* (CBPV) [Levy 2003] that we use as an intermediate language in Section 5. CBPV is a polarised λ -calculus that is structured around a fundamental dichotomy between *values* and *computations*. The primitives of CBPV provide logical mechanisms that control evaluation. In terms of expressivity, these mechanisms are strong enough to allow a faithful embedding of both CBN and CBV PCF in CBPV. The structure of CBPV ensures compatibility with computational effects such as printing, memory, etc. For the purposes of this paper, we will consider cost to be an effect: the embeddings of CBN and CBV into CBPV will explicitly mention a *command* that incurs evaluation cost whenever some should be incurred according to the source language operational semantics. This makes the embeddings essentially parametric in the evaluation costs: even if we decide to change the costs—e.g. by charging twice as much for a function call—we only need to slightly alter the embedding into CBPV and the accompanying cost-preservation proof. We thus demonstrate that CBPV is an extremely flexible and general intermediate language for recurrence extraction.

The extraction function and bounding relation for CBPV is defined in Section 6. It is here that we reap the benefits of the distinction between values and computations, for it is clear that they require different notions of bounding. This in turn explicates the differences in the bounding relations for our original call-by-value and call-by-name versions of PCF. The cost-preserving embeddings of the two languages that we give in Section 7 are fundamentally different: following Levy [2003], CBN types are translated to CBPV computation types, whereas CBV types are translated to CBPV value types.

We define a denotational semantics for the recurrence language PCF_c in Section 8 by introducing *sized domains*, which combine size and information orders. We observe that the standard model of PCF extends to a model in sized domains, in which the natural numbers are given their usual order. In Section 9, we use this model to give an end-to-end example of recurrence extraction for a non-structural recursion, in order to show that the recurrence we obtain indeed matches the one we would expect from an informal analysis. Finally, in Section 10 we sketch an extension of our work to a call-by-value source language with inductive types, following the approach of [Danner et al. 2015], and use it to extract a recurrence from a program implementing merge sort. Both this and the exponentiation example validate our claim that our technique formalizes the usual informal extract-and-solve approach to cost analysis.

2 THE SOURCE LANGUAGES: PCF

Our objective in this paper is to present a completely formal and provably correct way to extract recurrences from general-recursive functional programs. For the sake of clarity, we target a small formal calculus that illustrates the core features of this style of programming. Perhaps the most well-understood calculus of this sort is PCF [Plotkin 1977]. PCF essentially consists of a simply-typed λ -calculus, a base type of natural numbers, and some mechanism for obtaining fixed points. Almost every introductory book on programming language semantics includes a wealth of material on it: see e.g. [Gunter 1992; Mitchell 1996; Streicher 2006]. As discussed in the introduction, we would like to extract recurrences from both a CBN and a CBV variant of PCF.

The types, terms, and typing judgments of PCF are defined in Figure 3. The version we use comes with flat natural numbers, product, and function types. We write $\mathcal{T}_A^{\text{PCF}}$ for the (open and closed) terms of PCF of type A , and $\mathcal{V}_A^{\text{PCF}}$ for the CBV canonical forms of type A .

Types	A, B	$::=$	$\text{nat} \mid A \times B \mid A \rightarrow B$
Contexts	Γ	$::=$	$\cdot \mid \Gamma, x : A$
Numerical operations	op	$::=$	$+ \mid * \mid - \mid \div \mid \text{mod}$
Canonical forms (CBN)	V	$::=$	$\underline{n} \mid \langle M, N \rangle \mid \lambda x. M$
Canonical forms (CBV)	V, W, Z	$::=$	$\underline{n} \mid \langle V, W \rangle \mid \lambda x. M \mid \text{rec } f(x) = M$

$\frac{}{\Gamma, x : A, \Gamma' \vdash x : A}$	$\frac{}{\Gamma \vdash \underline{n} : \text{nat}}$	$\frac{\Gamma \vdash N : \text{nat} \quad \Gamma \vdash P, Q : A}{\Gamma \vdash \text{if } N \text{ then } P \text{ else } Q : A}$	$\frac{\Gamma \vdash M : \text{nat} \quad \Gamma \vdash N : \text{nat}}{\Gamma \vdash M \text{ op } N : \text{nat}}$
$\frac{\Gamma, x : A \vdash M : A}{\Gamma \vdash \text{fix } x. M : A}$	$\frac{\Gamma \vdash M : A_1 \quad \Gamma \vdash N : A_2}{\Gamma \vdash \langle M, N \rangle : A_1 \times A_2}$	$\frac{\Gamma \vdash P : A_1 \times A_2}{\Gamma \vdash \pi_i(P) : A_i}$	
$\frac{\Gamma, f : A \rightarrow B, x : A \vdash M : B}{\Gamma \vdash \text{rec } f(x) = M : A \rightarrow B}$	$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x. M : A \rightarrow B}$	$\frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash M N : B}$	

Fig. 3. PCF

Our natural numbers are introduced by constants $\underline{0}, \underline{1}, \dots$ and correspond to eager, rather than lazy, natural numbers. Instead of the usual predecessor and successor functions, we introduce five arithmetic operations on natural numbers, as well as the zero test if N then P else Q , which behaves like P if N is $\underline{0}$, and like Q otherwise. This choice of primitives gives them the flavor of *machine words*. We sometimes refer to these as “flat” natural numbers, as they admit a domain interpretation with a flat information order (bottom below all numerals, and no other relations).

The terms of the CBN and the CBV variants differ only on fixed points. In the case of CBN we may take fixed points at every type: if we have $M : A$ in terms of $x : A$, we may construct $\text{fix } x. M : A$. However, in CBV we are only allowed to take fixed points at function types, i.e. to make recursive function definitions: if we have $M : B$ in terms of an argument $x : A$ and a recursive variable $f : A \rightarrow B$, then we can obtain a recursively defined function $\text{rec } f(x) = M : A \rightarrow B$. This formulation of CBV PCF is commonly found in the literature; see e.g. [Levy 2006, §7], [Fiore 1996, §1.2.5], or [Plotkin and Power 2001].

The big-step semantics of the two variants of PCF are defined in Figure 4. Judgments of the form $M \Downarrow^n V$ are read as “ M evaluates to canonical form V incurring a cost n in the process.” Except the presence of the cost superscripts, these rules are standard in the literature on PCF, while the cost annotations follow Blelloch and Greiner [1995]. Each rule sums the costs incurred in the premises and adds additional cost for the evaluation step represented by that rule. For simplicity, we choose to charge one unit of cost for pair projections and (possibly recursive) function applications, and no cost otherwise. This could easily be adjusted to give any constant cost for each reduction, but—as recurrence extraction is syntax-directed—each reduction for the same term constructor must incur the same cost (e.g. we cannot charge 1 for reducing an if to the 0 branch and 2 for reducing it to the other branch). We have that

THEOREM 2.1 (DETERMINACY). *There is at most one pair n, V such that $M \Downarrow^n V$.*

We let $M \Downarrow \stackrel{\text{def}}{=} \exists n, V. M \Downarrow^n V$ which we pronounce as “ M is bounded,” and write $M \Uparrow$ to mean $\neg(M \Downarrow)$.

Rules for all variants

$$\frac{}{\underline{n} \downarrow^0 \underline{n}} \quad \frac{M \downarrow^a \underline{m} \quad N \downarrow^b \underline{n}}{M \text{ op } N \downarrow^{a+b} \underline{m \text{ op } n}} \quad \frac{}{\lambda x. M \downarrow^0 \lambda x. M}$$

$$\frac{N \downarrow^a \underline{0} \quad P \downarrow^b V}{\text{if } N \text{ then } P \text{ else } Q \downarrow^{a+b} V} \quad \frac{N \downarrow^a \underline{n+1} \quad Q \downarrow^b V}{\text{if } N \text{ then } P \text{ else } Q \downarrow^{a+b} V}$$

Call-by-name PCF

$$\frac{}{\langle M, N \rangle \downarrow^0 \langle M, N \rangle} \quad \frac{P \downarrow^a \langle M_1, M_2 \rangle \quad M_i \downarrow^b V_i}{\pi_i(P) \downarrow^{a+b+1} V_i} \quad \frac{M \downarrow^m \lambda x. P \quad P[N/x] \downarrow^n V}{MN \downarrow^{m+n+1} V} \quad \frac{M[\text{fix } x. M/x] \downarrow^n V}{\text{fix } x. M \downarrow^{n+1} V}$$

Call-by-value PCF

$$\frac{}{\text{rec } f(x) = P \downarrow^0 \text{rec } f(x) = P} \quad \frac{M \downarrow^a V \quad N \downarrow^b W}{\langle M, N \rangle \downarrow^{a+b} \langle V, W \rangle} \quad \frac{P \downarrow^n \langle V_1, V_2 \rangle}{\pi_i(P) \downarrow^{n+1} V_i}$$

$$\frac{M \downarrow^m \text{rec } f(x) = P \quad N \downarrow^n W \quad P[\text{rec } f(x) = P/f, W/x] \downarrow^k V}{MN \downarrow^{m+n+k+1} V} \quad \frac{M \downarrow^m \lambda x. P \quad N \downarrow^n W \quad P[W/x] \downarrow^k V}{MN \downarrow^{m+n+k+1} V}$$

Fig. 4. Big-step semantics for PCF

$$\widehat{0} \stackrel{\text{def}}{=} \mathbf{0} \quad \widehat{n+1} \stackrel{\text{def}}{=} \mathbf{1} \boxplus \widehat{n}$$

Types $A, B ::= \dots \mid \mathbb{C}$

Canonical forms (CBN only) $V ::= \dots \mid \widehat{n}$

Typing rules: The CBN rules of Figure 3, plus $\frac{\widehat{n} \in \{0, 1\}}{\Gamma \vdash \widehat{n} : \mathbb{C}}$ and $\frac{\Gamma \vdash M : \mathbb{C} \quad \Gamma \vdash N : \mathbb{C}}{\Gamma \vdash M \boxplus N : \mathbb{C}}$

Big-step semantics: The CBN rules of Figure 4, plus $\frac{\widehat{n} \in \{0, 1\}}{\widehat{n} \downarrow \widehat{n}}$ and $\frac{M \downarrow \widehat{m} \quad N \downarrow \widehat{n}}{M \boxplus N \downarrow \widehat{m+n}}$

Fig. 5. PCF with costs

3 THE RECURRENCE LANGUAGE: PCF WITH COSTS

The recurrences we aim to extract from PCF terms will themselves be expressed in an appropriate *recurrence language*, which will also be a version of PCF, and which we call *PCF with costs* (PCFc). As mentioned before, we would like the recurrence language to behave in a way that is as close to mathematics as possible. Hence, we choose PCFc to be a CBN language.

In addition to standard constructs, PCFc also sports an additional type of costs, which we denote by $\mathbb{C}$. We assume as little as possible about costs: there are constants for no cost and unit cost, denoted $\mathbf{0}$ and $\mathbf{1}$. Furthermore, PCFc comes with an operator that adds arbitrary costs together: if $M : \mathbb{C}$ and $N : \mathbb{C}$ we have $M \boxplus N : \mathbb{C}$. We define $\widehat{n}$ to be the right-associated sum $\mathbf{1} \boxplus \dots \boxplus \mathbf{0}$ of n unit costs. We write $\mathcal{T}_A^{\text{PCFc}}$ for the (open and closed) terms of PCFc of type A .

The operational semantics of PCFc consists of a judgment $M \downarrow V$, which is read as “term M evaluates to canonical form V .” Notice that this judgment does not record any evaluation costs, as we are not interested in the running time of recurrences. The rules for $M \downarrow V$ consist of all the CBN rules of Figure 4—with the cost annotations erased—along with two additional rules that handle the evaluation of costs: one which allows 0 and 1 to evaluate to themselves, and one which handles the summation of costs in the expected way. These facts are summarised in Figure 5.

However, the operational semantics of PCFc play a very limited rôle: they are only used to define the predicate $M \downarrow$ for PCFc. Instead, the central device used alongside PCFc is the *size order*. It consists of judgments of the form $\Gamma \vdash M \leq N : A$, which we read as “ M is bounded above by N at type A in context Γ .” The size order is used to compare the costs denoted by recurrences, as well as the sizes of values.

The full inductive definition of the size order is introduced in Figure 6. To explain its definition, we first need to discuss three auxiliary notions.

Syntactic Unfolding. The *n*th syntactic unfolding of a fixed point term, denoted by $\text{fix}_n x. M$, is in a sense the “*n*th approximation” to the fixed point $\text{fix } x. M$. Intuitively, $\text{fix}_n x. M$ is the term that runs for up to n recursive calls, and in the next call decides to diverge. It is defined by

$$\begin{aligned} \text{fix}_0 x. M &\stackrel{\text{def}}{=} \text{fix } x. x \\ \text{fix}_{n+1} x. M &\stackrel{\text{def}}{=} M[\text{fix}_n x. M/x] \end{aligned}$$

It is easy to see that

PROPOSITION 3.1. *If $\Gamma, x : A \vdash M : A$ then $\Gamma \vdash \text{fix}_n x. M : A$ for any $n \geq 0$.*

Monotone contexts. As is usual, PCF term contexts are defined in a manner similar to PCF terms. The difference is that PCF contexts have had one of their subterms replaced by a *hole*, denoted by $[]$. For example, the context $C \stackrel{\text{def}}{=} \lambda x. M([] + 3)$ is obtained by replacing the term N in $\lambda x. M(N + 3)$ by a hole $[]$. We may recover the original term $\lambda x. M(N + 3)$ by *filling* the hole with N , which we write as $C[N]$ and define in the usual way. Recall also that—unlike substitution—filling a hole may result in variable capture. For example, the variable x is no longer free in $C[x] \equiv \lambda x. M(x + 3)$.

Of the contexts, we choose certain ones to be *monotone*, which means that they preserve the size order $\leq$. In order to prove the bounding theorem about recurrence extraction into PCFc, we will need that the monotone contexts include at least the principal positions of elimination forms, as well as introduction forms for negative types (i.e. the bodies of functions and pairs), as stated in Figure 6. It is always permissible, and sometimes desirable, to make more contexts monotone (e.g. all of them). However, this places more requirements on semantic models of PCFc, so in this figure we only include those that are necessary for proving the bounding theorem. Monotone contexts are introduced by a typing judgment $C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B)$. The meaning of this judgment is this: the hole in C is assumed to represent a term $\Gamma \vdash M : A$, i.e. a term of type A in context Γ . When we fill the hole with a $\Gamma \vdash M : A$, the resultant term $C[M]$ has type B in context Δ . Note that the context may change, as C might bind some free variables.

PROPOSITION 3.2 (FILLING TYPING). $\Gamma \vdash M : A \wedge C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B) \implies \Delta \vdash C[M] : B$

Eliminative contexts. Of the monotone contexts, the (negative) *eliminative contexts* consist of a series of applications and projections. They are defined by

$$\mathcal{E} ::= [] \mid \pi_i(\mathcal{E}) \mid \mathcal{E}M$$

Eliminative contexts are strict in effects, and in particular they preserve divergence:

LEMMA 3.3 (PRESERVATION OF INFINITY). *For an eliminative context $\mathcal{E}$, $M \uparrow$ implies $\mathcal{E}[M] \uparrow$.*

$\frac{}{[] :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Gamma \vdash A)}$ $\frac{\Delta \vdash M : \mathbb{C} \quad C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \mathbb{C})}{M \boxplus C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \mathbb{C})}$ $\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta, x : B \vdash C)}{\lambda x. C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B \rightarrow C)}$ $\frac{\Delta \vdash M : B_1 \quad C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B_2)}{\langle M, C \rangle :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B_1 \times B_2)}$ $\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash A_1 \times A_2)}{\pi_i(C) :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash A_i)}$	$\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \text{nat}) \quad \Delta \vdash M, N : C}{\text{if } C \text{ then } M \text{ else } N :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash C)}$ $\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \mathbb{C}) \quad \Delta \vdash N : \mathbb{C}}{C \boxplus N :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \mathbb{C})}$ $\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B \rightarrow C) \quad \Delta \vdash N : B}{CN :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash C)}$ $\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B_1) \quad \Delta \vdash N : B_2}{\langle C, N \rangle :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B_1 \times B_2)}$ $\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \text{nat}) \quad \Delta \vdash N : \text{nat}}{C \text{ op } N :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \text{nat})}$
$\frac{C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \text{nat}) \quad \Delta \vdash N : \text{nat} \quad \text{op}^+ \in \{+, *\}}{M \text{ op}^+ C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash \text{nat})}$	
$\frac{\Gamma \vdash M : A}{\Gamma \vdash M \leq M : A} \text{ (refl)}$ $\frac{\Gamma \vdash M : \mathbb{C}}{\Gamma \vdash M \leq \mathbf{0} \boxplus M : \mathbb{C}} \text{ (zero)}$ $\frac{\Gamma, x : A \vdash M : B \quad \Gamma \vdash N : A}{\Gamma \vdash N[M/x] \leq (\lambda x. M)N : B} (\rightarrow_\beta)$ $\frac{\Gamma \vdash P : A \quad \Gamma \vdash Q : A}{\Gamma \vdash P \leq \text{if } \mathbf{0} \text{ then } P \text{ else } Q : A} (\text{if}_{\text{tt}})$ $\frac{}{\Gamma \vdash \underline{n} \text{ op } \underline{m} \leq \underline{n} \text{ op } \underline{m} : \text{nat}} \text{ (num}_\beta\text{)}$ $\frac{\Gamma \vdash M \leq N : A \quad C :: (\Gamma \vdash A) \rightsquigarrow_{\text{mon.}} (\Delta \vdash B)}{\Delta \vdash C[M] \leq C[N] : B} \text{ (ctx)}$ $\frac{\Gamma, x : A \vdash E : A \quad \Delta, z : A \vdash \mathcal{E}[z] : B \quad \forall n \geq 0. \Delta \vdash M \leq \mathcal{E}[\text{fix}_n x. E] : B}{\Delta \vdash M \leq \mathcal{E}[\text{fix } x. E] : B} \text{ (cpind)}$	$\frac{\Gamma \vdash M \leq N : A \quad \Gamma \vdash N \leq P : A}{\Gamma \vdash M \leq P : A} \text{ (trans)}$ $\frac{\Gamma \vdash M : \mathbb{C} \quad \Gamma \vdash N : \mathbb{C} \quad \Gamma \vdash P : \mathbb{C}}{\Gamma \vdash M \boxplus (N \boxplus P) \leq (M \boxplus N) \boxplus P : \mathbb{C}} \text{ (assoc)}$ $\frac{\Gamma \vdash M_1 : A_1 \quad \Gamma \vdash M_2 : A_2}{\Gamma \vdash M_i \leq \pi_i(\langle M_1, M_2 \rangle) : A_i} (\times_\beta)$ $\frac{\Gamma \vdash P : A \quad \Gamma \vdash Q : A}{\Gamma \vdash Q \leq \text{if } \underline{n+1} \text{ then } P \text{ else } Q : A} (\text{if}_{\text{ff}})$ $\frac{\Gamma \vdash N : \text{nat}}{\Gamma \vdash \underline{m} \bmod N \leq N - 1 : \text{nat}} \text{ (mod)}$ $\frac{\Gamma, x : A \vdash E : A}{\Gamma \vdash \text{fix}_{n+1} x. E \leq \text{fix}_n x. E : A} \text{ (rat)}$

Fig. 6. Monotone contexts and the size order for PCF_c

The rules of the size order. The rules of the size order given in Figure 6 consist of exactly those that are needed in order to prove the bounding theorem of §6. For example, we do not assume that $\underline{n} \leq \underline{n+1} : \text{nat}$ or $\mathbf{0} \leq \mathbf{1} : \mathbb{C}$, so that we can consider models where $\leq$ is actually equality, in addition to models where these inequalities do hold (such as the one given in Section 8).

The rules for $\leq$ can be naturally organised in three groups. The first group is the one that makes $\leq$ a preorder—i.e. (refl), (trans)—and preserved by monotone contexts (ctx). The second

group contains rules that encode small-step β -reduction: these ensure that $N \leq M$ whenever we would have had $M \rightarrow_{\beta} N$. This group contains the rules (**zero**) and (**assoc**), which ensure that $\widehat{n+m} \leq \widehat{n} + \widehat{m}$ for all $n, m \in \mathbb{N}$, as a simple induction shows.

Finally, the third and most interesting group consists of the two rules that concern fixed points. Both of these rules deal with sequences of terms that we call *rational chains*, i.e. sequences of the form $(\text{fix}_i x. M)_{i \in \omega}$, which semantically correspond to chains of the form $(f^i(\perp))_{i \in \omega}$. Rational chains are ubiquitous in syntactic and intensional models of PCF: see e.g. [Abramsky et al. 1996; Escardó and Ho 2009; Milner 1977; Pitts 1997].

The rule (**rat**) ensures that rational chains are *decreasing* in the size order, i.e. that

$$\dots \leq \text{fix}_2 x. M \leq \text{fix}_1 x. M \leq \text{fix}_0 x. M \stackrel{\text{def}}{=} \infty$$

Recall that the $(n+1)$ th syntactic unfolding may make more recursive calls than the n th, so it is a more defined term. Hence, this rule intuitively states that *a more defined recurrence is a tighter bound*. Note that as $\text{fix}_{n+1} x. M$ is defined to be $M[\text{fix}_n x. M/x]$, this rule is analogous to β for fix .¹

Finally, the rule (**cpind**) is a form of *computational induction* for PCFc. It ensures that this process of iterative tightening of bounds does not ‘overshoot’: if we can prove that $\text{fix}_n x. E$ is an upper bound for M for every $n \geq 0$, then so is $\text{fix } x. E$. Moreover, this rule ensures that computational induction can take place under any eliminative context $\mathcal{E}$.

We immediately obtain the following results.

LEMMA 3.4 (INFINITE LOOP IS A TOP ELEMENT (INFINITY)). $\Gamma \vdash M \leq \text{fix } x. x : A$ for all $\Gamma \vdash M : A$.

PROOF. Use the rule (**rat**), $n = 0$ for a fresh variable $x \notin \text{Vars}(\Gamma)$. □

LEMMA 3.5 (BOUNDED TERMS ARE A LOWER SET). $M \downarrow \wedge N \leq M : A \implies N \downarrow$

PROOF. By induction on $N \leq M : A$. □

4 RECURRENCE EXTRACTION FOR CALL-BY-NAME AND CALL-BY-VALUE

We now have enough details in place to show how to extract recurrences for both call-by-name and call-by-value PCF.

4.1 Call-by-Value

The extraction procedure for CBV is very similar to that of [Danner et al. 2015, 2013], and is displayed in Figure 7. To begin, we map each type A of CBV PCF to a type

$$\|A\| \stackrel{\text{def}}{=} \mathbb{C} \times \langle\!\langle A \rangle\!\rangle$$

of PCFc, where $\langle\!\langle A \rangle\!\rangle$ is defined by induction on A . We call this the type of *complexities* for A . A complexity for A consists of a pair: its first component is the *cost* of evaluating a term of type A to a value, and its second component is the *potential* of the resultant value. If $E : \|A\|$ we write $E_c \stackrel{\text{def}}{=} \pi_1(E) : \mathbb{C}$ and $E_p \stackrel{\text{def}}{=} \pi_2(E) : \langle\!\langle A \rangle\!\rangle$ to refer to these two components respectively.

The notion of potential is crucial in CBV extraction. There are varying interpretations about the nature of potentials: in one reading, they—directly or indirectly—encode information about the *size* of values; in another, they represent the future cost of using that value, to which we often refer as *use-cost*. The definition of $\langle\!\langle A \rangle\!\rangle$ is also given in Figure 7. Perhaps the most interesting clause is $\langle\!\langle A \rightarrow B \rangle\!\rangle \stackrel{\text{def}}{=} \langle\!\langle A \rangle\!\rangle \rightarrow \langle\!\langle B \rangle\!\rangle$. We may read this as follows. The cost of using a value of function

¹We do not require β for fix itself as a size order axiom in order to prove the bounding theorem, essentially because in the case of the bounding theorem for fixed points, we will apply a fixed point induction principle to instead reason about fix_n . The β axioms are generally used in the proof of the bounding theorem to head expand, showing that a redex is in the relation if its reduct is, but this particular reduction does not come up.

$$\begin{aligned}
\|x\| &\stackrel{\text{def}}{=} \langle 0, x \rangle \\
\|\underline{n}\| &\stackrel{\text{def}}{=} \langle 0, \underline{n} \rangle \\
\|M\{+, * \}N\| &\stackrel{\text{def}}{=} \langle \|M\|_c \boxplus \|N\|_c, \|M\|_p \{+, * \} \|N\|_p \rangle \\
\|M\{-, \div \}\underline{n}\| &\stackrel{\text{def}}{=} \langle \|M\|_c, \|M\|_p \{-, \div \}\underline{n} \rangle \\
\|M\{-, \div \}N\| &\stackrel{\text{def}}{=} \langle \|M\|_c \boxplus \|N\|_c, \|M\|_p \rangle \text{ (if } N \neq \underline{n}) \\
\|M \bmod N\| &\stackrel{\text{def}}{=} \langle \|M\|_c \boxplus \|N\|_c, \|N\|_p - 1 \rangle \\
\langle A_1 \times A_2 \rangle &\stackrel{\text{def}}{=} \langle A_1 \rangle \times \langle A_2 \rangle \\
\langle A \rightarrow B \rangle &\stackrel{\text{def}}{=} \langle A \rangle \rightarrow \|B\| \\
\langle \text{nat} \rangle &\stackrel{\text{def}}{=} \text{nat} \\
\| \text{if } N \text{ then } P \text{ else } Q \| &\stackrel{\text{def}}{=} \|N\|_c +_c \text{ if } \|N\|_p \text{ then } \|P\| \text{ else } \|Q\| \\
\langle M, N \rangle &\stackrel{\text{def}}{=} \langle \|M\|_c \boxplus \|N\|_c, \langle \|M\|_p, \|N\|_p \rangle \rangle \\
\|\pi_i(M)\| &\stackrel{\text{def}}{=} \langle 1 \boxplus \|M\|_c, \pi_i(\|M\|_p) \rangle \\
\|\lambda x. M\| &\stackrel{\text{def}}{=} \langle 0, \lambda x. \|M\| \rangle \\
\|MN\| &\stackrel{\text{def}}{=} 1 \boxplus \|M\|_c \boxplus \|N\|_c +_c (\|M\|_p \|N\|_p) \\
\|\text{rec } f(x) = M\| &\stackrel{\text{def}}{=} \langle 0, \text{fix } f. \lambda x. \|M\| \rangle
\end{aligned}$$

$$\text{If } \Gamma \vdash E : \|A\| \text{ and } \Gamma \vdash C : \mathbb{C} \text{ we write } \begin{cases} \Gamma \vdash E_c \stackrel{\text{def}}{=} \pi_1(E) : \mathbb{C} \\ \Gamma \vdash E_p \stackrel{\text{def}}{=} \pi_2(E) : \langle A \rangle \\ \Gamma \vdash C +_c E \stackrel{\text{def}}{=} \langle C \boxplus \pi_1(E), \pi_2(E) \rangle : \|A\| \end{cases}$$

Fig. 7. Recurrence extraction for CBV PCF

type—i.e. a λ -expression—is expressed as a function itself. This function that maps the size/use-cost of a value of type A to a pair of an evaluation cost (the cost of evaluating the application of that λ -expression to a value with that use-cost) and the use-cost of a value of type B . This is consistent with the idea that in CBV *variables represent values*.

Following that, we define a map that extracts a recurrence from each term of CBV PCF, which we also denote by $\|-\|$ (see Danner et al. [2015] for an explanation of most of the cases). Arithmetic operations $M \text{ op } N$ are chosen to have zero cost in addition to the cost of evaluating their inputs (since for big- O bounds we mainly count recursive calls), but for their potential we distinguish some cases, guided by two constraints needed by our proof of the bounding theorem: the potential must be an upper bound on the value of the operation, and must be monotone in the potentials of the operands wherever mathematically possible. Addition/multiplication are monotone in both M and N , so we use addition/multiplication to combine the potentials of M and N , which will be an upper bound on $M\{+, *\}N$. In general, subtraction and division are monotone in M but antimonotone in N , so we take the potential to be that of M , which is an upper bound on $M\{-, \div\}N$. However, for subtraction or division with N a numeric constant there will be no monotonicity obligation for N (because it does not vary), so we can perform the $\{-, \div\}$ in the potential. This precision is useful for analyzing algorithms that subtract/divide by a constant in recursive calls. Finally, mod is not monotone in either position, so we cannot use mod in the potential; but it is bounded by both M and $N - 1$, and we somewhat arbitrarily choose the later. For algorithms (e.g. GCD) that recur on $x\{-, \div, \text{mod}\}y$, we would need more sophisticated tracking of monotone and antimonotone

positions in the recurrence language. The other new case is for recursive functions, and interprets them using fixed points in PCFc.

Extending $\llbracket - \rrbracket$ to contexts pointwise, we have

THEOREM 4.1. *If $\Gamma \vdash M : A$ then $\llbracket \Gamma \rrbracket \vdash \llbracket M \rrbracket : \llbracket A \rrbracket$.*

It remains to state a theorem to the effect that each extracted recurrence $\llbracket M \rrbracket : \llbracket A \rrbracket$ encodes an upper bound for the evaluation of M . Because of the presence of function types, this is stated as a logical relation:

THEOREM 4.2 (CBV EXTRACTION). *There exist relations*

$$(M : \mathcal{T}_A^{PCF, closed}) \sqsubseteq_A (E : \mathcal{T}_{\llbracket A \rrbracket}^{PCFc, closed}) \quad \text{and} \quad (V : \mathcal{V}_A^{PCF, closed}) \sqsubseteq_A^{\text{val}} (E : \mathcal{T}_{\llbracket A \rrbracket}^{PCFc, closed})$$

such that

$$\begin{aligned} M \sqsubseteq_A E &\implies \text{if } E_c \downarrow \text{ then } \exists n, V. \begin{cases} M \downarrow^n V \\ \widehat{n} \leq E_c \\ V \sqsubseteq_A^{\text{val}} E_p \end{cases} \\ \underline{n} \sqsubseteq_{\text{nat}}^{\text{val}} E &\implies \underline{n} \leq E \\ \langle V_1, V_2 \rangle \sqsubseteq_{A_1 \times A_2}^{\text{val}} E &\implies \begin{cases} V_1 \sqsubseteq_{A_1}^{\text{val}} \pi_1(E) \\ V_2 \sqsubseteq_{A_2}^{\text{val}} \pi_2(E) \end{cases} \\ \lambda x. M \sqsubseteq_{A \rightarrow B}^{\text{val}} E &\implies \forall (N \sqsubseteq_A^{\text{val}} X). M[N/x] \sqsubseteq_B EX \\ \text{rec } f(x) = P \sqsubseteq_{A \rightarrow B}^{\text{val}} E &\implies \forall (N \sqsubseteq_A^{\text{val}} X). P[\text{rec } f(x) = P/f, N/x] \sqsubseteq_B EX \end{aligned}$$

and, moreover,

- (1) $V \sqsubseteq_A^{\text{val}} \llbracket V \rrbracket_p$ for any CBV PCF value $\cdot \vdash V : A$,
- (2) $M \sqsubseteq_A \llbracket M \rrbracket$ for any closed CBV PCF term $\cdot \vdash M : A$.

We prove this as a corollary of recurrence extraction for CBPV below. Intuitively, the *expression bounding relation* $M \sqsubseteq_A E$ says that the cost and value of a CBV PCF program M are predicted by E in the following sense: if the cost component of E terminates, then so does M , and the evaluation cost of M is bounded by the cost component of E according to the size order. Moreover, the value is bounded by the potential component of E , according to the value bounding relation $\sqsubseteq_A^{\text{val}}$. The value bounding relation is type-directed: at nat, it says that n is bounded by E according to the size order; for pairs, it says that the components are bounded; and for functions it says that the applications are bounded.

Relative to the bounding relation in [Danner et al. \[2015\]](#), the key change here for supporting general recursion involves the quantifiers in expression bounding. In that work, expression bounding was defined as “if $M \downarrow^n V$ then $\widehat{n} \leq E_c$ and $V \sqsubseteq_A^{\text{val}} E_p$ ” (if the source program evaluates, then the recurrence’s prediction is correct), though because all programs in the language considered there terminate, this is equivalent to “ $M \downarrow^n V$ and $\widehat{n} \leq E_c$ and $V \sqsubseteq_A^{\text{val}} E_p$.” Here, we assert this same guarantee *only if* the cost component of the recurrence itself terminates. We view a recurrence whose cost diverges as predicting an infinite cost for the program, so the expression bounding relation is vacuously true in this case, expressing that any program meets this bound.

4.2 Call-by-Name

The extraction procedure for CBN PCF is very different to the one for CBV. We will discover the precise reasons for that through CBPV in §7, but for now we will satisfy ourselves with the

$$\begin{aligned}
& \|x\| \stackrel{\text{def}}{=} x \\
& \|\underline{n}\| \stackrel{\text{def}}{=} \langle 0, \underline{n} \rangle \\
& \|\text{nat}\| \stackrel{\text{def}}{=} \mathbb{C} \times \text{nat} \\
& \|A_1 \times A_2\| \stackrel{\text{def}}{=} \|A_1\| \times \|A_2\| \\
& \|A \rightarrow B\| \stackrel{\text{def}}{=} \|A\| \rightarrow \|B\| \\
& \|M \text{ op } N\| \stackrel{\text{def}}{=} (\text{same as in CBV}) \\
& \|\text{if } N \text{ then } P \text{ else } Q\| \stackrel{\text{def}}{=} \|N\|_c +_A \text{ if } \|N\|_p \text{ then } \|P\| \text{ else } \|Q\| \\
& \|\langle M, N \rangle\| \stackrel{\text{def}}{=} \langle 1 +_{A_1} \|M\|, 1 +_{A_2} \|N\| \rangle \\
& \|\pi_i(M)\| \stackrel{\text{def}}{=} \pi_i(\|M\|) \\
& \|\lambda x. M\| \stackrel{\text{def}}{=} \lambda x. 1 +_B \|M\| \\
& \|MN\| \stackrel{\text{def}}{=} \|M\| \|N\| \\
& \|\text{fix } x. M\| \stackrel{\text{def}}{=} \text{fix } x. 1 +_A \|M\| \\
& c : \mathbb{C}, x : \mathbb{C} \times \text{nat} \vdash \alpha_{\text{nat}}(c, x) \stackrel{\text{def}}{=} \langle c \boxplus \pi_1(x), \pi_2(x) \rangle : \mathbb{C} \times \text{nat} \\
& c : \mathbb{C}, p : \|A_1\| \times \|A_2\| \vdash \alpha_{A_1 \times A_2}(c, p) \stackrel{\text{def}}{=} \langle \alpha_{A_1}(c, \pi_1(p)), \alpha_{A_2}(c, \pi_2(p)) \rangle : \|A_1\| \times \|A_2\| \\
& c : \mathbb{C}, f : \|A\| \rightarrow \|B\| \vdash \alpha_{A \rightarrow B}(c, f) \stackrel{\text{def}}{=} \lambda a. \alpha_B(c, f(a)) : \|A\| \rightarrow \|B\|
\end{aligned}$$

Fig. 8. Recurrence extraction for CBN PCF

following intuitions. When studying sundry versions of PCF we often speak of certain types as being *observable*, in the sense that the successful normalization of terms at those types provides some manifest information to the user. For example, every term $M : \text{nat}$ either diverges or is observed to converge to a constant $\underline{n}$. However, in CBN we are unable to observe a function: the only action we may perform is that of applying it to an argument. In contrast, in CBV *every type is observable*:² for example, when evaluating a term at function type we either expect divergence or convergence to a λ -expression.

In the case of CBN we only have one observable type, i.e. that of natural numbers. This is to say that we expect the end user to only evaluate terms of type nat . It follows that only complexities for those terms should contain cost components. Complexities at any other type will need to come with some mechanism for “pushing costs down to nat .” Thus, a complexity for a CBN type will not merely be a type of PCFc, but will also come with some kind of algebra structure that enables us to do that at every type.

Definition 4.3 (Cost algebra). A (PCFc) *cost algebra* $A = (A^\bullet, \alpha_A)$ consists of a type $A^\bullet$ and a PCFc term $c : \mathbb{C}, x : A^\bullet \vdash \alpha_A(c, x) : A^\bullet$ such that

$$c : \mathbb{C}, x : A^\bullet \vdash \alpha_A(c_1 \boxplus c_2, x) \leq \alpha_A(c_1, \alpha_A(c_2, x)) : A^\bullet$$

A PCFc cost algebra³ $A = (A^\bullet, \alpha_A)$ consists of a type $A^\bullet$, which we call its *carrier*, and a *structure map* $\alpha_A(c, x)$, which is expressed as a term in two free variables. The structure map allows one to “add costs” to any term of carrier type. From this point onwards we will abuse notation by not making a formal distinction between algebras $A = (A^\bullet, \alpha_A)$ and their carriers $A^\bullet$.

²There is a third paradigm which is close to CBN, but where termination at function type is observable. This is usually called the *lazy* paradigm; see [Abramsky 1990], [Riecke 1993], and [Levy 2003, §1.7.3].

³The analogous equation $x : A^\bullet \vdash x \leq \alpha_A(0, x)$ is true if $\leq$ additionally includes certain η -contractions, but this property is not used here so we omit these η rules for maximum generality.

The extraction procedure for CBN may be found in Figure 8. We begin by defining an algebra $(\|A\|, \alpha_A)$ for each type A of CBN PCF. In particular, nat is mapped to the “free algebra” $(\mathbb{C} \times \text{nat}, \alpha_{\text{nat}})$ on nat . For simplicity we use the notation $L +_A M \stackrel{\text{def}}{=} \alpha_A(L, M)$ whenever $L : \mathbb{C}$ and $M : \|A\|$.

We then inductively define a map $\|-\|$ from *typed* terms of PCF to terms of PCFc. This definition uses the aforementioned algebras, and so the output depends on the type of each input term.⁴ This definition uses the algebras to push the cost of evaluation towards the ground types. Indeed, costs here are not added at the elimination rules, but *at the introduction rules*. Consider the lazy product type $\text{nat} \times \text{nat}$ as an example. A pair $\langle M, N \rangle$ of this type is not directly observable, but contains two observable computations M and N . A computation that uses this pair costs one unit more than a computation that uses either M or N directly, because it must do the product projection.

This translation is well-typed:

THEOREM 4.4. *If $\Gamma \vdash M : A$ then $\|\Gamma\| \vdash \|M\| : \|A\|$.*

Furthermore, it satisfies a bounding theorem given as a logical relation:

THEOREM 4.5 (CBN EXTRACTION). *There exists a relation*

$$(M : \mathcal{T}_A^{\text{PCF}, \text{closed}}) \sqsubseteq (E : \mathcal{T}_{\|A\|}^{\text{PCFc}, \text{closed}})$$

such that

$$\begin{aligned} M \sqsubseteq_{\text{nat}} E &\implies \text{if } E_c \downarrow \text{ then } \exists n, V. \begin{cases} M \Downarrow^n V \\ \hat{n} \leq E_c : \mathbb{C} \\ V \leq E_p : \text{nat} \end{cases} \\ M \sqsubseteq_{A_1 \times A_2} E &\implies \begin{cases} \pi_1(M) \sqsubseteq_{A_1} \pi_1(E) \\ \pi_2(M) \sqsubseteq_{A_2} \pi_2(E) \end{cases} \\ M \sqsubseteq_{A \rightarrow B} E &\implies \forall (N \sqsubseteq_A X). MN \sqsubseteq_B EX \end{aligned}$$

and, moreover, $M \sqsubseteq_A \|M\|$ for any PCF term $\cdot \vdash M : A$.

Unlike in CBV, we do not distinguish between value and expression relations. The CBN relation for nat (or other observable types, if we had them) is analogous to the CBV expression relation, while the clauses for negative types simply say that the elimination forms preserve relatedness. Again, rather than proving this bounding theorem directly, we will obtain it as a corollary of a general theorem about CBPV.

5 THE INTERMEDIATE LANGUAGE: LEVY’S CALL-BY-PUSH-VALUE (CBPV)

We aim to prove both theorems of Section 4 at once. In order to do so, we will embed both CBN and CBV PCF in an *intermediate language*, and we will do so in a cost-preserving manner. Finally, we will define a recurrence extraction function for this intermediate language, and prove that it is correct. The composite of these two processes will prove the results of Section 4.

Our intermediate language of choice of Levy’s *Call-by-Push-Value* (CBPV) [Levy 2003]. CBPV is a polarised λ -calculus whose type structure provides ways of controlling evaluation. The typing judgments of the version of CBPV that we will use in this paper—which includes natural numbers and recursion—may be found in Figure 9.

CBPV is structured around a fundamental dichotomy between values and computations. In particular, it comes with two sorts of types: *value types*, usually denoted by A , and *computation*

⁴Formally, the function is defined on typing derivations/intrinsically typed syntax $\|\Gamma \vdash M : A\|$, but we elide the annotations for brevity.

Value Types	A	$::=$	$\text{nat} \mid A_1 \times A_2 \mid \underline{UB}$
Computation Types	$\underline{B}$	$::=$	$FA \mid \underline{B}_1 \& \underline{B}_2 \mid A \rightarrow \underline{B}$
Contexts	Γ	$::=$	$\cdot \mid \Gamma, x : A$
Numerical operations	op	$::=$	$+ \mid - \mid * \mid \div \mid \text{mod}$
Terminal Computations	T, S	$:=$	$\text{return } V \mid \langle M, N \rangle \mid \lambda x. M$

$\frac{}{\Gamma, x : A, \Gamma' \vdash_v x : A} \text{(var)}$	$\frac{}{\Gamma \vdash_v \tilde{n} : \text{nat}} \text{(nat}_{\tilde{n}})$
$\frac{\Gamma \vdash_c M : \underline{B}}{\Gamma \vdash_c \text{charge}. M : \underline{B}} \text{(ch)}$	$\frac{\Gamma, x : \underline{UB} \vdash_c M : \underline{B}}{\Gamma \vdash_c \text{fix } x. M : \underline{B}} \text{(fix)}$
$\frac{\Gamma \vdash_v M, N : \text{nat} \quad \Gamma, z : \text{nat} \vdash_c Q : \underline{B}}{\Gamma \vdash_c \text{calc } z \Leftarrow M \text{ op } N \text{ in } Q : \underline{B}} \text{(natop)}$	$\frac{\Gamma \vdash_v N : \text{nat} \quad \Gamma \vdash_c P : \underline{B} \quad \Gamma \vdash_c Q : \underline{B}}{\Gamma \vdash_c \text{ifz}(N; P; Q) : \underline{B}} \text{(natifz)}$
$\frac{\Gamma \vdash_v V_1 : A_1 \quad \Gamma \vdash_v V_2 : A_2}{\Gamma \vdash_v (V_1, V_2) : A_1 \times A_2} (\times I)$	$\frac{\Gamma \vdash_v V : A_1 \times A_2 \quad \Gamma, x_1 : A_1, x_2 : A_2 \vdash_c N : \underline{B}}{\Gamma \vdash_c \text{split}(V; (x_1, x_2). N) : \underline{B}} (\times E)$
$\frac{\Gamma \vdash_c M : \underline{B}}{\Gamma \vdash_v \text{thunk } M : \underline{UB}} \text{(UI)}$	$\frac{\Gamma \vdash_v V : \underline{UB}}{\Gamma \vdash_c \text{force } V : \underline{B}} \text{(UE)}$
$\frac{\Gamma \vdash_v V : A}{\Gamma \vdash_c \text{return } V : FA} \text{(FI)}$	$\frac{\Gamma \vdash_c M : FA \quad \Gamma, x : A \vdash_c N : \underline{B}}{\Gamma \vdash_c \text{bind } x \Leftarrow M \text{ in } N : \underline{B}} \text{(FE)}$
$\frac{\Gamma \vdash_c M : \underline{B}_1 \quad \Gamma \vdash_c N : \underline{B}_2}{\Gamma \vdash_c \langle M, N \rangle : \underline{B}_1 \& \underline{B}_2} (\& I)$	$\frac{\Gamma \vdash_c M : \underline{B}_1 \& \underline{B}_2}{\Gamma \vdash_c \pi_i(M) : \underline{B}_i} (\& E)$
$\frac{\Gamma, x : A \vdash_c M : \underline{B}}{\Gamma \vdash_c \lambda x. M : A \rightarrow \underline{B}} (\rightarrow I)$	$\frac{\Gamma \vdash_c M : A \rightarrow \underline{B} \quad \Gamma \vdash_v N : A}{\Gamma \vdash_c MN : \underline{B}} (\rightarrow E)$

Fig. 9. Call-by-Push-Value (CBPV) with natural numbers and recursion

types, usually denoted by $\underline{B}$. Value types often include certain ground types, *positive* products, and *thunks*. Computation types include *free* computation types, *negative* product types, and function types. Like in CBV, variables in CBPV may only denote values, and are thus assigned value types.

The terms at value types are exactly values: the introduction rules encode all the possible ways of constructing them. For example, if $V : A_1$ and $W : A_2$, we have $(V, W) : A_1 \times A_2$. However, we may *not* project the components of that pair back into value types. But if we have a term $N : \underline{B}$ at some computation type with two free variables $x_1 : A_1$ and $x_2 : A_2$, then the elimination rule allows us to split a pair of values into N , yielding $\text{split}((V, W); (x_1, x_2). N) : \underline{B}$. It is in this way that the invariant of value types is maintained.

As their name suggests, the terms at computation type may demonstrate some non-trivial computational behaviour, including effects (e.g. printing, output, etc). In the words of Paul Levy, “a value is, a computation does.” The most basic such type is the *free* (computation) type FA , where A is a value type. A term $M : FA$ is a *returner*, which may engage in some effectful behaviour before returning some value by normalizing to the form $\text{return } V$ where $V : A$. Given $M : FA$ and as $x : A \vdash N : \underline{B}$ we may form $\text{bind } x \Leftarrow M \text{ in } N : \underline{B}$. This term will first run M , effecting some changes and returning a value which will then be substituted for x . Computation types also include *negative products*, which we write as $\underline{B}_1 \& \underline{B}_2$, and which come with the usual projections. Finally,

$$\begin{array}{c}
\frac{}{\lambda x. M \Downarrow^0 \lambda x. M} \quad \frac{}{\text{return } V \Downarrow^0 \text{return } V} \quad \frac{}{\langle M, N \rangle \Downarrow^0 \langle M, N \rangle} \\
\frac{M \Downarrow^m T}{\text{force } (\text{thunk } M) \Downarrow^m T} \quad \frac{N[V_1/x_1, V_2/x_2] \Downarrow^m T}{\text{split}((V_1, V_2); (x_1, x_2). N) \Downarrow^m T} \quad \frac{M[\text{thunk } (\text{fix } x. M)/x] \Downarrow^n T}{\text{fix } x. M \Downarrow^n T} \\
\frac{P \Downarrow^m T}{\text{ifz}(\widetilde{0}; P; Q) \Downarrow^m T} \quad \frac{Q \Downarrow^m T}{\text{ifz}(\widetilde{n+1}; P; Q) \Downarrow^m T} \quad \frac{P[\widetilde{m \text{ op } n}/z] \Downarrow^m T}{\text{calc } z \Leftarrow \widetilde{m \text{ op } n} \text{ in } P \Downarrow^m T} \\
\frac{M \Downarrow^n T}{\text{charge}. M \Downarrow^{n+1} T} \quad \frac{M \Downarrow^m \text{return } V \quad N[V/x] \Downarrow^n T}{\text{bind } x \Leftarrow M \text{ in } N \Downarrow^{m+n} T} \\
\frac{P \Downarrow^m \langle M_1, M_2 \rangle \quad M_i \Downarrow^n T}{\pi_i(P) \Downarrow^{m+n} T} \quad \frac{M \Downarrow^m \lambda x. P \quad P[V/x] \Downarrow^n T}{MV \Downarrow^{m+n} T}
\end{array}$$

Fig. 10. Big-step semantics for CBPV

they include function types: given a computation $M : \underline{B}$ in a free value variable $x : A$, we may form $\lambda x. M : A \rightarrow \underline{B}$. Notice that $A \rightarrow \underline{B}$ is a ‘mixed type,’ which mirrors the fact that variables may only stand for values.

Only one thing remains, and that is to overcome the restriction of variables to value types. This is achieved by introducing a value type $U\underline{B}$, the so-called type of *thunks* of $\underline{B}$, for every computation type $\underline{B}$. Each computation $M : \underline{B}$ may be *thunked* as $\text{thunk } M : U\underline{B}$, and each $N : U\underline{B}$ may be *forced* into a computation $\text{force } M : \underline{B}$. Thunk types are crucial in simulating CBN. They are also necessary in the CBPV fixed point rule, which assumes that the recursive call is given in the form of a thunk.

Our version of CBPV also comes with ‘flat’ natural numbers as a value type. As in PCF, these are introduced by an infinite number of constants $\widetilde{0}, \widetilde{1}, \dots$. We may use natural numbers in either of two ways. First, using the term $\text{ifz}(\widetilde{n}; P; Q)$ we may query them on whether they are zero, following the computation P if so, and Q otherwise. Second, using the construct $\text{calc } v \Leftarrow \widetilde{m \text{ op } n} \text{ in } N$ we may *calculate* one of four numerical operations on them, and substitute the result into some variable $v : \text{nat}$ that is free in the ‘continuation’ N .

Finally, we equip our version of CBPV with a single effect, namely the one that incurs a unit cost: for each computation $M : \underline{B}$ we may form a computation $\text{charge}. M : \underline{B}$. In section 7 we will use this effect in our embedding of CBN and CBV into CBPV to record the places where the $+1$ ’s appear in the operational semantics of PCF.

Figure 10 introduces big-step semantics for CBPV. The judgment $M \Downarrow^n T$ means that the closed CBPV term M evaluates to the *terminal computation* T , incurring a cost n in the process. Terminal computations include $\text{return } V$, pairs, and λ -expressions. Unlike the operational semantics of PCF, costs are incurred here only when they are explicitly mentioned as effects: with the exception of the rule for $\text{charge}. M$, all the other rules sum the costs of their premises.

6 RECURRENCE EXTRACTION FOR CBPV & THE BOUNDING THEOREM

Having introduced our intermediate language, we will now show how to extract recurrences from its terms. Furthermore, we will prove the central result of this paper, the *Bounding Theorem*, which ensures that extracted recurrences express upper bounds for the running times of CBPV terms.

Compared with the results of [Danner et al. 2015], our results are a significant technical improvement. First, the distinction between *potentials* and *complexities* is now formally motivated: our

techniques demonstrate that *values have potentials*, whereas *computations have complexities*. We thus get a very clear conceptual basis for higher-order recurrences. Second, the fact that our version of CBPV comes with an explicit cost effect makes this result reusable and extensible: our theorem applies not just to a specific source language, but to any source language that can be faithfully embedded in CBPV. All we need to do is encode its terms in CBPV in a way that explicitly records where evaluation costs are incurred, and then invoke the bounding theorem.

6.1 CBPV Recurrence Extraction

The extraction procedure for CBPV incorporates elements of the extraction procedures for both CBV and CBN given in §4. It is rather close to the monad/algebra categorical semantics of CBPV given in [Levy 2006, §3] in the special case of the writer monad ($\mathbb{C} \times -$).

To begin, we translate CBPV types to PCFc types as follows: a value type A is translated to a PCFc type $\langle\langle A \rangle\rangle$, its so-called type of *potentials*, and a computation type $\underline{B}$ is translated to a *cost algebra* $\|\underline{B}\| \stackrel{\text{def}}{=} (\|\underline{B}\|^\bullet, \alpha_{\underline{B}})$. Again, the idea is that values only have potential (future, indirect use-cost), whereas computations may be evaluated now, hence incurring (immediate, direct) costs. Because of the interdependence between value and computation types introduced by $F(-)$ and $U(-)$, the translation—which may be found in Figure 11—is defined in a mutually recursive manner. For the sake of precision we have maintained the formal distinction between algebras $\|\underline{B}\|$ and their carriers $\|\underline{B}\|^\bullet$. As before, we abbreviate $\alpha_{\underline{B}}(L, E)$ by $L +_{\underline{B}} E$.

Next, we give a translation of CBPV terms to PCFc terms. In Figure 11 we introduce the judgments $\Gamma \vdash_v V \searrow P : A$ and $\Gamma \vdash_c M \searrow Q : \underline{B}$. The first one is read as “the CBPV term V of value type A translates to the PCFc term P in context Γ .” The interpretation of the second judgment is similar, but involves terms of computation type. We have that

THEOREM 6.1 (CBPV-TO-PCFc TRANSLATION).

- (1) If $\Gamma \vdash_v V : A$ then there is a unique PCFc term P such that $\Gamma \vdash_v V \searrow P : A$, and $\langle\langle \Gamma \rangle\rangle \vdash P : \langle\langle A \rangle\rangle$.
- (2) If $\Gamma \vdash_c M : \underline{B}$ then there is a unique PCFc term Q such that $\Gamma \vdash_c M \searrow Q : \underline{B}$, and $\langle\langle \Gamma \rangle\rangle \vdash Q : \|\underline{B}\|^\bullet$.

Hence, whenever $\Gamma \vdash_v V \searrow P : A$ we know that this P is unique. We write $\langle\langle V \rangle\rangle$ for it, so that $\langle\langle \Gamma \rangle\rangle \vdash \langle\langle V \rangle\rangle : \langle\langle A \rangle\rangle$. We also write $\|\underline{M}\|$ for the unique Q such that $\Gamma \vdash_c M \searrow Q : \underline{B}$. In that case we have $\langle\langle \Gamma \rangle\rangle \vdash \|\underline{M}\| : \|\underline{B}\|^\bullet$. From this point onwards we again confuse algebras with their carriers.

The main ideas of this recurrence extraction are: implementing the effect charge M in CBPV as adding unit cost using the algebra on the computation type $\underline{B}$ of M , and undoing some of the finer operational distinctions that are made in CBPV but not in PCFc. For example, both eager and lazy products in CBPV are interpreted as products in PCFc, and both split and projections in CBPV are interpreted as projections in PCFc. Because of the latter, it may be surprising that we can obtain a bounding theorem that works for the CBV features of CBPV. A precedent for this kind of translation of CBV into a CBN-like setting is the adequacy theorem of [Plotkin and Power 2001] for CBV PCF with algebraic effects: the main requirement in that result is a strong monad with a *strict* strength, and the writer monad $\mathbb{C} \times -$ is such a monad.

6.2 The Bounding Theorem

We now wish to argue that the extraction procedure for CBPV is correct via a logical relation that generalizes the CBV and CBN ones. We define two relations

$$(V : \mathcal{T}_A^{\text{CBPV, closed}}) \lesssim_A^{\text{val}} (E : \mathcal{T}_{\langle\langle A \rangle\rangle}^{\text{PCFc, closed}}) \quad (M : \mathcal{T}_{\underline{B}}^{\text{CBPV, closed}}) \lesssim_{\underline{B}}^c (E : \mathcal{T}_{\|\underline{B}\|}^{\text{PCFc, closed}})$$

by induction on types; the full definitions are shown in Figure 12. $V \lesssim_A^{\text{val}} E$ is read as “the potential of the value $V : A$ is bounded above by E ,” and $M \lesssim_{\underline{B}}^c E$ is read as “the complexity of the computation

$$\begin{array}{l}
\langle\langle \text{nat} \rangle\rangle \stackrel{\text{def}}{=} \text{nat} \qquad \qquad \qquad \|FA\| \stackrel{\text{def}}{=} (\mathbb{C} \times \langle\langle A \rangle\rangle, \alpha_{FA}) \\
\langle\langle A_1 \times A_2 \rangle\rangle \stackrel{\text{def}}{=} \langle\langle A_1 \rangle\rangle \times \langle\langle A_2 \rangle\rangle \qquad \|A \rightarrow B\| \stackrel{\text{def}}{=} (\langle\langle A \rangle\rangle \rightarrow \|B\|^\bullet, \alpha_{A \rightarrow B}) \\
\langle\langle UB \rangle\rangle \stackrel{\text{def}}{=} \|B\|^\bullet \qquad \|B_1 \& B_2\| \stackrel{\text{def}}{=} (\|B_1\|^\bullet \times \|B_2\|^\bullet, \alpha_{B_1 \& B_2}) \\
\\
c : \mathbb{C}, x : \mathbb{C} \times \langle\langle A \rangle\rangle \vdash \alpha_{FA}(c, x) \stackrel{\text{def}}{=} \langle c \boxplus \pi_1(x), \pi_2(x) \rangle : \mathbb{C} \times \langle\langle A \rangle\rangle \\
c : \mathbb{C}, f : \langle\langle A \rangle\rangle \rightarrow \|B\|^\bullet \vdash \alpha_{A \rightarrow B}(c, f) \stackrel{\text{def}}{=} \lambda a. \alpha_B(c, f(a)) : \langle\langle A \rangle\rangle \rightarrow \|B\|^\bullet \\
c : \mathbb{C}, p : \|B_1\|^\bullet \times \|B_2\|^\bullet \vdash \alpha_{B_1 \& B_2}(c, p) \stackrel{\text{def}}{=} \langle \alpha_{B_1}(c, \pi_1(p)), \alpha_{B_2}(c, \pi_2(p)) \rangle : \|B_1\|^\bullet \times \|B_2\|^\bullet \\
\\
\textbf{Note:} \text{ whenever } \Gamma \vdash E : \|FA\|, \text{ we let } \begin{cases} \Gamma \vdash E_c \stackrel{\text{def}}{=} \pi_1(E) : \mathbb{C} \\ \Gamma \vdash E_p \stackrel{\text{def}}{=} \pi_2(E) : \langle\langle A \rangle\rangle \end{cases} \\
\\
\frac{}{\Gamma, x : A, \Gamma' \vdash_v x \searrow x : A} \text{(var)} \qquad \frac{}{\Gamma \vdash_v \tilde{n} \searrow \underline{n} : \text{nat}} \qquad \frac{\Gamma, x : UB \vdash_c M \searrow P : B}{\Gamma \vdash_c \text{fix } x. M \searrow \text{fix } x. P : B} \\
\\
\frac{\Gamma \vdash_v V_1 \searrow W_1 : A_1 \quad \Gamma \vdash_v V_2 \searrow W_2 : A_2}{\Gamma \vdash_v (V_1, V_2) \searrow \langle W_1, W_2 \rangle : A_1 \times A_2} \qquad \frac{\Gamma \vdash_v V \searrow P : A_1 \times A_2 \quad \Gamma, x_1 : A_1, x_2 : A_2 \vdash_c N \searrow Q : B}{\Gamma \vdash_c \text{split}(V; (x_1, x_2). N) \searrow Q[\pi_1(P)/x_1, \pi_2(P)/x_2] : B} \\
\\
\frac{\Gamma \vdash_v V \searrow P : A}{\Gamma \vdash_c \text{return } V \searrow \langle 0, P \rangle : FA} \qquad \frac{\Gamma \vdash_c M \searrow Q : FA \quad \Gamma, x : A \vdash_c N \searrow R : B}{\Gamma \vdash_c \text{bind } x \leftarrow M \text{ in } N \searrow Q_c +_B R[Q_p/x] : B} \\
\\
\frac{\Gamma, x : A \vdash_c M \searrow P : B}{\Gamma \vdash_c \lambda x. M \searrow \lambda x. P : A \rightarrow B} \qquad \frac{\Gamma \vdash_c M \searrow P : A \rightarrow B \quad \Gamma \vdash_v N \searrow Q : A}{\Gamma \vdash_c MN \searrow PQ : B} \\
\\
\frac{\Gamma \vdash_c M \searrow P : B_1 \quad \Gamma \vdash_c N \searrow Q : B_2}{\Gamma \vdash_c \langle M, N \rangle \searrow \langle P, Q \rangle : B_1 \& B_2} \qquad \frac{\Gamma \vdash_c M \searrow P : B_1 \& B_2}{\Gamma \vdash_c \pi_i(M) \searrow \pi_i(P) : B_i} \\
\\
\frac{\Gamma \vdash_c M \searrow P : B}{\Gamma \vdash_v \text{thunk } M \searrow P : UB} \qquad \frac{\Gamma \vdash_v V \searrow P : UB}{\Gamma \vdash_c \text{force } V \searrow P : B} \\
\\
\frac{\Gamma \vdash_c M \searrow P : B}{\Gamma \vdash_c \text{charge}. M \searrow 1 +_B P : B} \qquad \frac{\Gamma \vdash_v V \searrow W : \text{nat} \quad \Gamma \vdash_c M \searrow P : B \quad \Gamma \vdash_c N \searrow Q : B}{\Gamma \vdash_v \text{ifz}(V; M; N) \searrow \text{if } W \text{ then } P \text{ else } Q : B} \\
\\
\frac{\Gamma \vdash_v V \searrow M : \text{nat} \quad \Gamma \vdash_v W \searrow N : \text{nat} \quad \Gamma, z : \text{nat} \vdash_c P \searrow Q : B \quad \mathbf{op}^+ \in \{+, *\}}{\Gamma \vdash_c \text{calc } z \leftarrow V \mathbf{op}^+ W \text{ in } P \searrow Q[M \mathbf{op}^+ N/z] : B} \\
\\
\frac{\Gamma \vdash_v V \searrow M : \text{nat} \quad \Gamma, z : \text{nat} \vdash_c P \searrow Q : B \quad \mathbf{op}^- \in \{-, \div\}}{\Gamma \vdash_c \text{calc } z \leftarrow V \mathbf{op}^- \tilde{n} \text{ in } P \searrow Q[M \mathbf{op}^- \tilde{n}/z] : B} \\
\\
\frac{\Gamma \vdash_v V \searrow M : \text{nat} \quad \Gamma \vdash_v W \searrow N : \text{nat} \quad \Gamma, z : \text{nat} \vdash_c P \searrow Q : B \quad \mathbf{op}^- \in \{-, \div\} \quad N \neq \tilde{n}}{\Gamma \vdash_c \text{calc } z \leftarrow V \mathbf{op}^- W \text{ in } P \searrow Q[M/z] : B} \\
\\
\frac{\Gamma \vdash_v V \searrow M : \text{nat} \quad \Gamma \vdash_v W \searrow N : \text{nat} \quad \Gamma, z : \text{nat} \vdash_c P \searrow Q : B}{\Gamma \vdash_c \text{calc } z \leftarrow V \bmod W \text{ in } P \searrow Q[N - 1/z] : B}
\end{array}$$

Fig. 11. Recurrence extraction for CBPV

$$\begin{array}{ll}
\tilde{n} \lesssim_{\text{nat}}^{\text{val}} E \stackrel{\text{def}}{=} \underline{n} \leq E & M \lesssim_{FA}^c E \stackrel{\text{def}}{=} E_c \downarrow \implies \exists n, V. \begin{cases} M \Downarrow^n \text{ return } V \\ \tilde{n} \leq E_c \\ V \lesssim_A^{\text{val}} E_p \end{cases} \\
(V_1, V_2) \lesssim_{A_1 \times A_2}^{\text{val}} E \stackrel{\text{def}}{=} \begin{cases} V_1 \lesssim_{A_1}^{\text{val}} \pi_1(E) \\ V_2 \lesssim_{A_2}^{\text{val}} \pi_2(E) \end{cases} & M \lesssim_{A \rightarrow B}^c E \stackrel{\text{def}}{=} \forall (N \lesssim_A^{\text{val}} X). M N \lesssim_B^c E X \\
\text{thunk } M \lesssim_{UB}^{\text{val}} E \stackrel{\text{def}}{=} M \lesssim_B^c E & M \lesssim_{B_1 \& B_2}^c E \stackrel{\text{def}}{=} \begin{cases} \pi_1(M) \lesssim_{B_1}^c \pi_1(E) \\ \pi_2(M) \lesssim_{B_2}^c \pi_2(E) \end{cases}
\end{array}$$

Fig. 12. Bounding relations for CBPV

$M : \underline{B}$ is bounded above by E .” The value relation is similar to the CBV value bounding relation, with the clause for U -types corresponding to the switch from value bounding to expression bounding in function types. The computation relation is similar to the CBN bounding relation, with the clause for F -types corresponding to the CBV expression relation.

We then establish that this relation holds at the diagonal, i.e. that

THEOREM 6.2 (BOUNDING THEOREM).

- (1) If $\cdot \vdash V : A$ then $V \lesssim_A^{\text{val}} \llbracket V \rrbracket$.
- (2) If $\cdot \vdash M : \underline{B}$ then $M \lesssim_B^c \llbracket M \rrbracket$.

While technical, the proof of this theorem is unsurprising, and may be found in an extended version of this paper. Amongst other things, it requires an auxiliary notion of *costless weak head reduction*. It is denoted by $M \longrightarrow_0 N$, and encodes all weak head reductions, plus a ‘commuting conversion’ that makes negative eliminators commute with the unit cost construct, e.g. $\pi_i(\text{charge}. M) \longrightarrow_0 \text{charge}. \pi_i(M)$. The main lemmas used are:

LEMMA 6.3 (ALGEBRA MONOTONICITY). $E \leq E' : \mathbb{C} \implies E +_B M \leq E' +_B M$

LEMMA 6.4 (BOUND WEAKENING).

- (1) $V \lesssim_A^{\text{val}} E \wedge E \leq E' : \llbracket A \rrbracket \implies V \lesssim_A^{\text{val}} E'$
- (2) $M \lesssim_B^c E \wedge E \leq E' : \llbracket B \rrbracket \implies M \lesssim_B^c E'$

LEMMA 6.5 (HEAD EXPANSION/REDUCTION).

- (1) $M' \lesssim_B^c E \wedge M \longrightarrow_0 M' \implies M \lesssim_B^c E$
- (2) $M \lesssim_B^c E \wedge M \longrightarrow_0 M' \implies M' \lesssim_B^c E$

LEMMA 6.6 (UNIT CHARGE). $M \lesssim_B^c E \implies \text{charge}. M \lesssim_B^c 1 +_B E$

THEOREM 6.7 (COMPACTNESS). $\mathcal{E}[\text{fix } x. M] \downarrow \implies \exists n \geq 0. \mathcal{E}[\text{fix}_n x. M] \downarrow$

LEMMA 6.8 (BOUNDS FOR RECURSION).

- (1) (*Infinity*) $E \uparrow \implies M \lesssim_B^c E$.
- (2) (*Infinity-Algebra*) $L \uparrow \implies M \lesssim_B^c L +_B E$.
- (3) (*Fixed Point Induction*) $\left(\forall n \geq 0. M \lesssim_B^c \text{fix}_n x. E \right) \implies M \lesssim_B^c \text{fix } x. E$

Compactness is known to hold for PCF. It can be obtained through a standard denotational semantics—see e.g. Pitts [1994, §10.9]—or through syntactic methods—see [Pitts 1997, §4].

A (PCF type, CBN)	$A^\dagger$ (CBPV+ computation type)
nat	$F(\text{nat})$
$A_1 \times A_2$	$A^\dagger \& B^\dagger$
$A \rightarrow B$	$U(A^\dagger) \rightarrow B^\dagger$

$x_1 : A_1, \dots, x_n : A_n \vdash M : A$	$x_1 : U(A_1^\dagger), \dots, x_n : U(A_n^\dagger) \vdash_c M^\dagger : A^\dagger$
x	force x
$\underline{n}$	return $\tilde{n}$
$M \text{ op } N$	bind $m \Leftarrow M^\dagger$ in bind $n \Leftarrow N^\dagger$ in calc $v \Leftarrow m \text{ op } n$ in return v
if N then P else Q	bind $n \Leftarrow N^\dagger$ in ifz(n ; $P^\dagger$; $Q^\dagger$)
$\langle M, N \rangle$	$\langle \text{charge}. M^\dagger, \text{charge}. N^\dagger \rangle$
$\pi_i(M)$	$\pi_i(M^\dagger)$
$\lambda x. M$	$\lambda x. (\text{charge}. M^\dagger)$
MN	$M^\dagger(\text{thunk } N^\dagger)$
fix $x. M$	fix $x. (\text{charge}. M^\dagger)$

Fig. 13. Call-by-name PCF to CBPV translation

7 EMBEDDING PCF INTO CBPV

The only piece of the puzzle that remains is to show how both CBN and CBV PCF can be embedded in CBPV in a cost-preserving way. At a high level, we follow [Levy 2003, §2.7] in defining a term translation, but insert charge expressions wherever the operational semantics of CBV or CBN PCF incurs a cost. We then prove that this translation preserves and reflects the operational semantics, including the cost.

7.1 CBN

In order to embed CBN into CBPV, we must recall that

- (1) the only *observable* type of CBN PCF is nat, and
- (2) the variables of CBN terms represent *thunks*, as we may substitute them with arbitrary terms.

These two facts directly lead us to a translation of any CBN type A to a CBPV *computation* type $A^\dagger$. As nat is observable, it is mapped to $F(\text{nat})$ (it returns a value). Products are compositionally mapped to *negative* products. Finally, as variables stand for thunks, we must thunk the domain when translating the function type.

The translation of both types and terms is defined in Figure 13. Notice that computation steps where explicit evaluation to canonical form is necessary—e.g. for testing whether a number is zero—are punctuated by the appearance of the bind $x \Leftarrow (-)$ in $(-)$ construct. Furthermore, notice that a variable represents a thunk, so it must be forced. Finally, notice the appearance of charge. $(-)$ whenever the operational semantics of PCF would impose a unit charge.

Extending $(-)^{\dagger}$ to contexts pointwise, we have that

PROPOSITION 7.1. *If $\Gamma \vdash M : A$ in call-by-name PCF then $\Gamma^\dagger \vdash_c M^\dagger : A^\dagger$.*

To prove that the operational semantics are preserved and reflected under this translation we follow the technique of [Levy 2006, §7]. We inductively define a relation $\Rightarrow_{\text{cbn}}$ between terms of PCF and terms of CBPV. The relation includes one rule for each line of the table of Figure 13, e.g.

$$\frac{M \Rightarrow_{\text{cbn}} M' \quad N \Rightarrow_{\text{cbn}} N'}{MN \Rightarrow_{\text{cbn}} M'(\text{thunk } N')}$$

Moreover, the definition also includes the rule

$$\frac{M \Rightarrow_{\text{cbn}} M'}{M \Rightarrow_{\text{cbn}} \text{force thunk } M'}$$

which—for technical reasons—inserts enough occurrences of `force thunk` ($-$). Evidently,

PROPOSITION 7.2. $M \Rightarrow_{\text{cbn}} M^\dagger$

We can then show that

THEOREM 7.3 (CBN BISIMULATION).

- (1) If $M \Downarrow^n V$ and $M \Rightarrow_{\text{cbn}} M'$, then $M' \Downarrow^n T'$ for some T' with $V \Rightarrow_{\text{cbn}} T'$.
- (2) If $M' \Downarrow^n T'$ and $M \Rightarrow_{\text{cbn}} M'$, then $M \Downarrow^n V$ for some V with $V \Rightarrow_{\text{cbn}} T'$.

PROOF. We first show that $M \Rightarrow_{\text{cbn}} M'$ and $N \Rightarrow_{\text{cbn}} N'$ implies $M[N/x] \Rightarrow_{\text{cbn}} M'[\text{thunk } N'/x]$, by an easy induction on $M \Rightarrow_{\text{cbn}} M'$. We then show the main claims. In the first case, we proceed by induction on $M \Downarrow^n V$, followed by local inductions on $M \Rightarrow_{\text{cbn}} M'$. In the second case, we proceed by induction on $M' \Downarrow^n T'$, followed by local inductions on $M \Rightarrow_{\text{cbn}} M'$. $\square$

COROLLARY 7.4 (CBN PRESERVATION).

- (1) If $M \Downarrow^n V$ then $M^\dagger \Downarrow^n T$ with $V \Rightarrow_{\text{cbn}} T$.
- (2) If $M^\dagger \Downarrow^n T$ then $M \Downarrow^n V$ with $V \Rightarrow_{\text{cbn}} T$.

7.2 CBV

In order to embed CBV into CBPV, we must recall that

- (1) in CBV, termination at every type is observable, and returns a *value*
- (2) the variables of CBV terms represent *values*

In short: *all terms of CBV are returners*. Thus, we aim to translate each type CBV type A to a CBPV value type A^* , and each term $x : A \vdash M : B$ to a term $x : A^* \vdash_c M^* : F(B^*)$ of CBPV.

Translating types is then fairly evident: `nat` is translated to itself, and products are compositionally mapped to *positive* products. Finally, $A \rightarrow B$ is mapped to $A^* \rightarrow F(B^*)$, which is then thunked in order to become a value type.

On the level of terms there are two translations: one takes CBV PCF terms to CBPV terms of $F(-)$ type, and the other one takes CBV PCF values to CBPV values. They are both defined in Figure 14. Notice once again that computation steps where explicit evaluation to canonical form is necessary—e.g. for function application, or construction of pairs—necessitate a use of the bind $x \leftarrow (-)$ in $(-)$ construct, which forces evaluation. Furthermore, notice that a variable represents a value, so it must be returned. Finally, notice the appearance of charge. $(-)$ whenever the operational semantics of PCF would impose a unit charge.

Extending $(-)^*$ to contexts pointwise, we have that

PROPOSITION 7.5.

- (1) If $\Gamma \vdash V : A$ is a PCF value, then $\Gamma^* \vdash_v V^{\text{val}} : A^*$.
- (2) If $\Gamma \vdash M : A$ in call-by-value PCF, then $\Gamma^* \vdash_c M^* : F(A^*)$.

It is even easier than CBN to establish that the operational semantics of CBV are preserved and reflected under this translation.

THEOREM 7.6 (CBV PRESERVATION).

- (1) If $M \Downarrow^n V$ then $M^* \Downarrow^n V^{\text{val}}$.
- (2) If $M^* \Downarrow^n \text{return } W$ then there exists V such that $M \Downarrow^n V$ and $W \equiv_\alpha V^{\text{val}}$.

A (PCF type, CBV)		A* (CBPV+ value type)	
nat		nat	
$A_1 \times A_2$		$A^* \times B^*$	
$A_1 \rightarrow A_2$		$U(A^* \rightarrow F(B^*))$	

$x_1 : A_1, \dots, x_n : A_n \vdash M : A$	$x_1 : A_1^*, \dots, x_n : A_n^* \vdash_c M^* : F(A^*)$
x	return x
$\frac{n}{M \text{ op } N}$	return $\tilde{n}$
if N then P else Q	bind $m \Leftarrow M^*$ in bind $n \Leftarrow N^*$ in calc $v \Leftarrow m \text{ op } n$ in return v
$\langle M, N \rangle$	bind $n \Leftarrow N^*$ in ifz(n ; P^* ; Q^*)
$\pi_i(M)$	bind $x \Leftarrow M^*$ in bind $y \Leftarrow N^*$ in return (x, y)
$\lambda x. M$	bind $p \Leftarrow M^*$ in charge. split(p ; (x_1, x_2)). return x_i
MN	return thunk $\lambda x. M^*$
rec $f(x) = M$	bind $f \Leftarrow M^*$ in bind $x \Leftarrow N^*$ in charge. (force f) x
	return thunk (fix $f. \lambda x. M^*$)

$x_1 : A_1, \dots, x_n : A_n \vdash V : A$	$x_1 : A_1^*, \dots, x_n : A_n^* \vdash_v V^{\text{val}} : A^*$
x	x
$\frac{n}{\lambda x. M}$	$\tilde{n}$
rec $f(x) = M$	thunk $\lambda x. M^*$
	thunk (fix $f. \lambda x. M^*$)

Fig. 14. Call-by-value PCF to CBPV translation

PROOF. First, we need to establish two lemmas:

- $V^* \equiv_{\alpha} \text{return } V^{\text{val}}$ for a call-by-value PCF value V
- $(M[V/x])^* \equiv_{\alpha} M^*[V^{\text{val}}/x]$

These follow straightforwardly by induction on values and terms respectively. The first claim then follows by induction on $M \Downarrow^n V$, and the second by induction on $M^* \Downarrow^n \text{return } W$. $\square$

7.3 Completing the Circle

Combining the bounding theorem of Section 6 and these results, we can prove the main theorems of §4. For Theorem 4.2 we will make use of head reduction, viz. Lemma 6.5. We then define

$$M \sqsubset_A^{\text{val}} E \stackrel{\text{def}}{=} V^{\text{val}} \lesssim_{A^*}^{\text{val}} E \qquad M \sqsubset_A E \stackrel{\text{def}}{=} M^* \lesssim_{F(A^*)}^c E$$

and calculate that the desired conditions hold. We then do the same for Theorem 4.5 and

$$M \sqsubset_A E \stackrel{\text{def}}{=} M^{\dagger} \lesssim_{A^{\dagger}}^c E$$

8 DENOTATIONAL SEMANTICS FOR PCF WITH COSTS

This section is concerned with the development of a denotational semantics for PCFc. Having completed the step of recurrence extraction, we would now like to interpret our *syntactic recurrences*—expressed in the formal language of PCF—into *semantic recurrences*. These should be as close as possible to good old-fashioned recurrences found in textbooks on algorithms.

Because of the presence of recursion, this endeavour requires a modicum of *domain theory*. Standard references to domain theory include [Abramsky and Jung 1994; Stoltenberg-Hansen et al. 1994], whereas references on the denotational semantics of PCF in domains may be found in §2. We assume knowledge of the standard definition of *complete partial order (cpo)*, i.e. a partial order

$(D, \sqsubseteq)$ with least upper bounds of directed subsets, and of Scott-continuous functions. We also recall the notion of an ω -chain $(x_i)_{i \in \omega}$ in a cpo, i.e. an increasing sequence of elements $x_1 \sqsubseteq x_2 \sqsubseteq \dots$.

8.1 Sized Domains

The relation usually modelled by the denotational semantics of PCF is that of *equality*: the semantics of PCF terms is *sound*, in that it is stable under evaluation. However, our main notion here is that of the axiomatic size order, and in order to model it we introduce the following kind of domain.

Definition 8.1. A *sized domain* (with joins) $(D, \sqsubseteq, \leq, \vee, \perp, \mathbf{0}, 1)$ consists of

- a set D ,
- a partial order $\sqsubseteq$, which we call the *information order* on D , and
- a preorder $\leq$, which we call the *size order* on D , and
- an operation $\vee : D \times D \rightarrow D$

such that

- (1) $(D, \sqsubseteq)$ is a cpo with least element $\perp$,
- (2) $(D, \leq)$ has chosen binary *joins* (maximums/least upper bounds) given by $\vee : D \times D \rightarrow D$ and least element $\mathbf{0}$,
- (3) $\vee : D \times D \rightarrow D$ is continuous with respect to $(D, \sqsubseteq)$,
- (4) $x \sqsubseteq y$ implies $y \leq x$, and
- (5) if $(x_i)_{i \in \omega}$ is a ω -chain in D and $\forall i \in \omega. z \leq x_i$, then $z \leq \bigsqcup_{i \in \omega} x_i$.

We immediately know that

PROPOSITION 8.2. $\perp$ is the greatest element (up to iso) with respect to $(D, \leq)$.

PROOF. For any $x \in D$ we have that $\perp \sqsubseteq x$, and hence $x \leq \perp$ by axiom (4). Thus, $\perp$ is a greatest element in $(D, \leq)$. However, as $(D, \leq)$ is not a partial order, it is only unique up to isomorphism. $\square$

The most interesting axioms in this definition are (4) and (5). The first one formalizes the idea that *a more defined bound is a better bound*. Consider the function $f = \{1 \mapsto 2, x \mapsto \perp\}$ on the flat domain of natural numbers. When interpreted as a recurrence, f gives us an upper bound for the input 1, but no information about any other input. In fact, Proposition 8.2 tells us that $\perp$ is the greatest element in the size order: it represents the *infinite value*. Thus, the recurrence f maps all elements other than 1 to the trivial upper bound of infinity. The recurrence $g = \{1 \mapsto 2, 2 \mapsto 2, x \mapsto \perp\}$ is more well-defined than f , as it is also defined (i.e. not $\perp$) at input 2: we have $f \sqsubseteq g$. As $\perp$ represents infinity, this is indeed a *tighter* bound, and axiom (4) stipulates that we should be able to treat it as such: we should have $g \leq f$. Axiom (4) is closely related to rule (rat) of the axiomatic size order of Figure 6, and is indeed used to verify its soundness.

Axiom (5) expresses the idea that *a recursively defined recurrence is an upper bound if all its approximants are*. Recall that fixed points in PCF are interpreted as the least upper bound of a chain $(f^i(\perp))_{i \in \omega}$ that is increasing in the information order $\sqsubseteq$. Intuitively, each element of this chain is more well-defined than its predecessor. Axiom (5) requires that if every extra step of definition $f^i(\perp)$ of a recurrence is indeed an upper bound for an element, then so must be the limit $\bigsqcup_{i \in \omega} x_i$. Axiom (5) is very close to the rule (cpind), and is used to verify its soundness.

We define the category **SizeDom** to consist of sized domains and Scott-continuous functions. The morphisms need not preserve any aspect of the size order; for example, we do not require *size monotonicity* in general, because only the PCF_c monotone contexts C are required to be interpreted as size-monotonic functions. It is also easy to see that this category is cartesian closed, as it inherits

products and exponentials from the underlying category $\mathbf{Cpo}$ of complete partial orders and Scott-continuous functions: we only need to define the size order and the chosen joins—which are in both cases defined pointwise—and show that the axioms of Definition 8.1 follow.

8.2 Interpreting PCF with Costs into Sized Domains

Defining a semantic interpretation of PCF into sized domains is reasonably straightforward. To each type A of PCF we associate a sized domain $\llbracket A \rrbracket = (D_A, \sqsubseteq_A, \leq_A, \vee_A, \perp_A, 0_A, 1_A)$. We interpret both the natural numbers nat and the type $\mathbb{C}$ of costs by the *flat sized domain of natural numbers*, which is defined to be $\mathbb{N}_\perp \stackrel{\text{def}}{=} (\mathbb{N} \cup \{\perp\}, \sqsubseteq, \leq, \vee, \perp, 0, 1)$ where $\perp \notin \mathbb{N}$. Its underlying set consists of the natural numbers augmented with a ‘fresh’ element $\perp$, which simultaneously models non-termination and infinity. Writing $\leq_{\mathbb{N}}$ for the usual order on natural numbers, we define

$$x \sqsubseteq y \stackrel{\text{def}}{=} x = \perp \text{ or } x = y, \quad x \leq y \stackrel{\text{def}}{=} y = \perp \text{ or } x \leq_{\mathbb{N}} y, \quad x \vee y \stackrel{\text{def}}{=} \begin{cases} \perp & \text{if } x = \perp \text{ or } y = \perp \\ \max_{\mathbb{N}} \{x, y\} & \text{otherwise} \end{cases}$$

It is easy to see that the requisite axioms are satisfied. The rest of the types are interpreted by

$$\llbracket A_1 \times A_2 \rrbracket \stackrel{\text{def}}{=} \llbracket A_1 \rrbracket \times \llbracket A_2 \rrbracket \quad \llbracket A \rightarrow B \rrbracket \stackrel{\text{def}}{=} \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket$$

where $[X \rightarrow Y]$ denotes the set of Scott-continuous functions from a cpo $(X, \sqsubseteq_X)$ to a cpo $(Y, \sqsubseteq_Y)$, and where both the information and size orders are given pointwise.

To each judgment $\Gamma \vdash M : A$, where $\Gamma \equiv x_1 : A_1, \dots, x_n : A_n$, we associate a function $\llbracket \Gamma \vdash M : A \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$ where $\llbracket \Gamma \rrbracket \stackrel{\text{def}}{=} \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket$. This definition is mostly standard—see e.g. [Streicher 2006, §3.2]. However, there are two unexpected elements, which arise because the PCF monotone contexts C must be interpreted as size-monotone functions in order to interpret the syntactic size order axiom (ctx), and because we interpret nat with the usual $0 \leq 1 \leq 2 \dots \leq \infty$ size order to facilitate reasoning about semantic recurrences. First, because if C then P else Q is a monotone context, we must “monotonize” the conditional. We do this by sending ∞ to ∞ , and using the chosen joins to make the interpretation for ≥ 1 dominate⁵ the answer for 0:

$$\llbracket \Gamma \vdash \text{if } N \text{ then } P \text{ else } Q : A \rrbracket (\vec{d}) \stackrel{\text{def}}{=} \begin{cases} \perp_A & \text{if } \llbracket \Gamma \vdash N : \text{nat} \rrbracket (\vec{d}) = \perp \\ \llbracket P \rrbracket (\vec{d}) & \text{if } \llbracket \Gamma \vdash N : \text{nat} \rrbracket (\vec{d}) = 0 \\ \llbracket P \rrbracket (\vec{d}) \vee_A \llbracket Q \rrbracket (\vec{d}) & \text{otherwise} \end{cases}$$

Second, because the operands of an arithmetic operation are monotone contexts, we must interpret the operations by a monotone upper bound of their actual numerical value. While addition and multiplication are already monotone, for subtraction, division and modulo, we use the following definitions, where $\mathbf{op}^- \in \{-, \div\}$:⁶

$$\begin{aligned} \llbracket \Gamma \vdash M \mathbf{op}^- n : \text{nat} \rrbracket (\vec{d}) &\stackrel{\text{def}}{=} \llbracket \Gamma \vdash M : \text{nat} \rrbracket (\vec{d}) \mathbf{op}^- n \\ \llbracket \Gamma \vdash M \mathbf{op}^- N : \text{nat} \rrbracket (\vec{d}) &\stackrel{\text{def}}{=} \llbracket \Gamma \vdash M : \text{nat} \rrbracket (\vec{d}) \\ \llbracket \Gamma \vdash M \bmod N : \text{nat} \rrbracket (\vec{d}) &\stackrel{\text{def}}{=} \llbracket \Gamma \vdash N : \text{nat} \rrbracket (\vec{d}) - 1 \end{aligned}$$

⁵The standard definition of the conditional along with $0 \leq 1$ and (ctx) would imply $\llbracket P \rrbracket = \llbracket \text{if } 0 \text{ then } P \text{ else } Q \rrbracket \leq \llbracket \text{if } 1 \text{ then } P \text{ else } Q \rrbracket = \llbracket Q \rrbracket$ for arbitrary terms P and Q .

⁶We could in fact remove $\{-, \div\}$ with non-numerals and mod from PCF, because they are never used as the target of the recurrence extraction in Section 6. In general, there is a choice about whether to perform this monotoneization in recurrence extraction or in the semantic interpretation. The above recurrence extraction from CBPV to PCF already does monotoneization, eliminating all mods and $\{-, \div\}$ with non-numerals. However, we prefer to leave these operations in PCF and interpret them similarly here to also support the other style of recurrence extraction into PCF.

$$\begin{aligned}
\text{rec exp}(n) &= \text{if } n \text{ then } 1 \text{ else let } z \leftarrow \text{exp}(n/2), y \leftarrow \text{if } n \bmod 2 \text{ then } 1 \text{ else } 2 \text{ in } z * z * y \\
\llbracket \text{exp} \rrbracket &= \langle 0, \text{fix exp. } \lambda n. \text{if } n \text{ then } \langle 0, \underline{1} \rangle \\
&\quad \text{else } \langle \widehat{3} \boxplus (\text{exp}(n/2))_c, \\
&\quad (\text{exp}(n/2))_p * (\text{exp}(n/2))_p * (\text{if } \underline{1} \text{ then } \underline{1} \text{ else } \underline{2}) \rangle \rangle
\end{aligned}$$

Fig. 15. Divide-and-conquer exponentiation and its extracted recurrence

The ‘new’ clauses interpret $\llbracket 0 \rrbracket(\vec{d}) \stackrel{\text{def}}{=} 0$, $\llbracket 1 \rrbracket(\vec{d}) \stackrel{\text{def}}{=} 1$, and

$$\llbracket \Gamma \vdash M \boxplus N : \mathbb{C} \rrbracket(\vec{d}) \stackrel{\text{def}}{=} \begin{cases} \perp & \text{if } \llbracket M \rrbracket(\vec{d}) = \perp \text{ or } \llbracket N \rrbracket(\vec{d}) = \perp \\ \llbracket M \rrbracket(\vec{d}) +_{\mathbb{N}} \llbracket N \rrbracket(\vec{d}) & \text{otherwise} \end{cases}$$

We have that

THEOREM 8.3 (SOUNDNESS). $\Gamma \vdash M \leq N : A \implies \forall \vec{d} \in \llbracket \Gamma \rrbracket. \llbracket M \rrbracket(\vec{d}) \leq_A \llbracket N \rrbracket(\vec{d})$

Moreover, we can use standard logical relations techniques to show that this interpretation is *adequate* in an appropriately relaxed sense.

THEOREM 8.4 (ADEQUACY).

- (1) If $M : \text{nat}$, then $\llbracket M \rrbracket = m$ implies $M \downarrow \underline{k}$ for some $k \leq_{\mathbb{N}} m$.
- (2) If $M : \mathbb{C}$, then $\llbracket M \rrbracket = m$ implies $M \downarrow \widehat{k}$ for some $k \leq_{\mathbb{N}} m$.

We anticipate the existence of other models of PCFc in sized domains, which we expect to obtain by varying the orders on the interpretation of nat , or even $\times$ and $\rightarrow$. For example, there should be a more precise discrete model, where all size orders are interpreted by equality, all cost bounds are exact, and the potentials include the programs themselves.

9 EXAMPLE: EXPONENTIATION

Next, we show how the above technique works on an example of a non-structurally-recursive function, which is not supported by Danner et al. [2015]. Consider the recursive specification for “exponentiation-by-squaring”:

$$2^n = \begin{cases} (2^{n/2})^2 & \text{if } n \text{ is even} \\ (2^{(n-1)/2})^2 * 2 & \text{otherwise} \end{cases}$$

We implement this in CBV PCF in Figure 15, using let expressions as syntactic sugar for β -expansions and recalling that we are performing floor division so that $n/2 = (n-1)/2$ when n is odd. We extract a recurrence in PCFc using the rules of Figure 7 and show the result in the same figure. In order to keep the term from becoming unwieldy, we assume a few additional size order axioms on PCFc terms beyond those of Figure 6, all of which are true in the semantics of sized domains (we write equality for the size order $\leq$ in both directions):

- the opposite direction of rules (**assoc**) and (**zero**), so that $(\boxplus, 0)$ forms a monoid,
- the opposite direction of the β rules, as well as η equality rules for $\times$ and $\rightarrow$ types,
- the ‘commuting conversion’ $\pi_1(\text{if } N \text{ then } M_1 \text{ else } M_2) = \text{if } N \text{ then } \pi_1(M_1) \text{ else } \pi_1(M_2)$, and
- if $\underline{k}$ then N else $N = N$

It is $\|\exp\|_p$ that is the desired recurrence, and we observe that if $T(n) = \pi_1(\|\exp\|_p n)$, then

$$T(0) = 0 \qquad T(n) = 3 + T(n/2)$$

which is the expected recurrence from an informal analysis.

When n is even, the potential is an over-approximation of the actual value: we get $2*((\exp(n/2))_p)^2$ whether n is even or odd because the recurrence extraction sends $n \bmod 2$ to 1 in order to be monotone. We leave a more precise treatment of non-monotone operations to future work.

10 INDUCTIVE TYPES

Our previous work [Danner et al. 2015] only permitted call-by-value functional programs with strictly positive inductive datatypes and structural recursion. The present paper extends these techniques to general non-structural recursion, which often leads to significantly more concise and efficient code. It is thus interesting to examine the combination of this presentation with the datatypes supported by the previous techniques of Danner et al. [2015]. We believe that this combination is achievable. In this section we illustrate it by a careful treatment of the case of lists in CBV PCF, and we briefly discuss the general case. In future work we plan to extend our techniques to *general recursive types*, which subsume lazy/CBN PCF lists as well as other coinductive types, though we expect that defining the bounding relation will be somewhat more challenging.

Figure 16 introduces the necessary additions to the syntax of CBV PCF and CBPV, and also extends the cost-preserving embedding, recurrence extraction and bounding relation to the new constructs. Intuitively, we choose here to represent a list by its length. This is useful for analyzing many algorithms whose running time does not depend on the list elements; if the cost does depend on list elements, one should instead represent a list by its length and maximum element, or some other more precise information. We could perform this abstraction of lists as lengths in the semantics, as in Danner et al. [2015], but this would introduce quite a bit of notational overhead. For the sake of simplicity, we instead translate a list to its length in the recurrence extraction phase, defining the type of potentials of lists to be nat . This way we do not need to add lists to PCFc.

The only essentially non-trivial clause is the recurrence extraction for case $(V; M_{\text{nil}}; (x, xs). M_{\text{cons}})$. Intuitively, this should map to a PCFc zero test if W then P else Q , where V , M_{nil} , and M_{cons} extract to W , P , and Q respectively. However, Q has two free variables, x and xs , which must be substituted for. These correspond to the potentials of the head and tail of V . The answer for xs is immediate, as its potential is bounded above by $W - 1$. As we have taken the potential of a list to be its length, we do not have any information about the size of the list elements. Hence, we choose the potential of x to be ∞ (fix $x. x$). This gives an upper bound in all cases, and is useful when the cost of the function to be analyzed does not depend on the size of the list elements: in that case the ∞ substituted for x will drop out of the recurrence at some point, so we will still obtain a finite bound. To complete the development, we extend the proof of the bounding theorem by adding a straightforward case for each new term construct. To do so we need the size order axioms $M \leq (M + 1) - 1$, $0 \leq 1$ and fix $x. x \leq \mathcal{E}[\text{fix } x. x]$ for eliminative contexts $\mathcal{E}$, all of which are valid in the standard semantics.

Though we have not written out all of the details, we expect that we can extend the above to allow potentials other than list length (e.g. length and maximum element), and to all strictly positive inductive types. Briefly, we need to add the corresponding inductive types to PCFc, add new size order axioms, extract case expressions as case expressions, and then modify the notion of sized domains to provide an interpretation of lists and other inductive types, as in Danner et al. [2015]. By doing so, the syntactic recurrence becomes a cost-annotated version of the original CBPV program that preserves a maximal amount of information about the inductive type values in the original program. Different sized domains can then be used to give different abstractions of size for inductive types, e.g. length for lists, height for (labelled) trees, or more complicated metrics,

CBV PCF

Types $A ::= \dots \mid \text{list } A$ **Canonical forms (CBV)** $V, W ::= \dots \mid \text{nil} \mid \text{cons}(V, W)$

$$\begin{array}{c}
 \frac{}{\Gamma \vdash \text{nil} : \text{list } A} \quad \frac{\Gamma \vdash M : A \quad \Gamma \vdash N : \text{list } A}{\Gamma \vdash \text{cons}(M, N) : \text{list } A} \quad \frac{\Gamma \vdash N : \text{list } A \quad \Gamma \vdash M_{\text{nil}} : B \quad \Gamma, x : A, xs : \text{list } A \vdash M_{\text{cons}} : B}{\Gamma \vdash \text{case } n \text{ of } \{ \text{nil} \mapsto M_{\text{nil}} \mid \text{cons}(x, xs) \mapsto M_{\text{cons}} \} : B} \\
 \\
 \frac{}{\text{nil} \downarrow^0 \text{nil}} \quad \frac{N \downarrow^n \text{nil} \quad M_{\text{nil}} \downarrow^m Z}{\text{case } N \text{ of } \{ \text{nil} \mapsto M_{\text{nil}} \mid \text{cons}(x, xs) \mapsto M_{\text{cons}} \} \downarrow^{n+m} Z} \\
 \\
 \frac{M \downarrow^m V \quad N \downarrow^n W}{\text{cons}(M, N) \downarrow^{m+n} \text{cons}(V, W)} \quad \frac{N \downarrow^n \text{cons}(V, W) \quad M_{\text{cons}}[V/x, W/xs] \downarrow^m Z}{\text{case } N \text{ of } \{ \text{nil} \mapsto M_{\text{nil}} \mid \text{cons}(x, xs) \mapsto M_{\text{cons}} \} \downarrow^{n+m} Z}
 \end{array}$$

CBPV

Value types $A ::= \dots \mid \text{list } A$

$$\begin{array}{c}
 \frac{}{\Gamma \vdash_{\text{v}} \text{nil} : \text{list } A} \quad \frac{\Gamma \vdash_{\text{v}} M : A \quad \Gamma \vdash_{\text{v}} N : \text{list } A}{\Gamma \vdash_{\text{v}} \text{cons}(M, N) : \text{list } A} \\
 \\
 \frac{\Gamma \vdash_{\text{v}} N : \text{list } A \quad \Gamma \vdash_{\text{c}} M_{\text{nil}} : \underline{B} \quad \Gamma, x : A, xs : \text{list } A \vdash_{\text{c}} M_{\text{cons}} : \underline{B}}{\Gamma \vdash_{\text{c}} \text{case}(N; M_{\text{nil}}; (x, xs). M_{\text{cons}}) : \underline{B}} \\
 \\
 \frac{M_{\text{nil}} \Downarrow^n T}{\text{case}(\text{nil}; M_{\text{nil}}; (x, xs). M_{\text{cons}}) \Downarrow^n T} \quad \frac{M_{\text{cons}}[V/x, W/xs] \Downarrow^n T}{\text{case}(\text{cons}(V, W); M_{\text{nil}}; (x, xs). M_{\text{cons}}) \Downarrow^n T}
 \end{array}$$

CBV PCF translation to CBPV

A (PCF type, CBV)	A^* (CBPV+ value type)
$\text{list } A$	$\text{list } A^*$
$x_1 : A_1, \dots, x_n : A_n \vdash M : A$	$x_1 : A_1^*, \dots, x_n : A_n^* \vdash_{\text{c}} M^* : F(A^*)$
nil	return nil
$\text{cons}(M, N)$	$\text{bind } m \leftarrow M^* \text{ in bind } n \leftarrow N^* \text{ in return cons}(m, n)$
$\text{case } M \text{ of } \{ \text{nil} \mapsto M_{\text{nil}} \mid \text{cons}(x, xs) \mapsto M_{\text{cons}} \}$	$\text{bind } m \leftarrow M^* \text{ in case } (m; M_{\text{nil}}^*; (x, xs). M_{\text{cons}}^*)$

CBPV recurrence extraction

$$\begin{array}{c}
 \langle\langle \text{list } A \rangle\rangle \stackrel{\text{def}}{=} \text{nat} \quad \frac{}{\Gamma \vdash_{\text{v}} \text{nil} \searrow \underline{0} : \text{list } A} \quad \frac{\Gamma \vdash_{\text{v}} W \searrow P : \text{list } A}{\Gamma \vdash_{\text{v}} \text{cons}(V, W) \searrow P + \underline{1} : \text{list } A} \\
 \\
 \frac{\Gamma \vdash_{\text{v}} V \searrow W : \text{list } A \quad \Gamma \vdash_{\text{c}} M_{\text{nil}} \searrow P : \underline{B} \quad \Gamma, x : A, xs : \text{list } A \vdash_{\text{c}} M_{\text{cons}} \searrow Q : \underline{B}}{\Gamma \vdash_{\text{c}} \text{case}(V; M_{\text{nil}}; (x, xs). M_{\text{cons}}) \searrow \text{if } W \text{ then } P \text{ else } Q[\text{fix } x. x/x, W - \underline{1}/xs] : \underline{B}} \\
 \\
 \text{nil} \lesssim_{\text{list } A}^{\text{val}} E \stackrel{\text{def}}{=} \underline{0} \leq E \quad \text{cons}(V, W) \lesssim_{\text{list } A}^{\text{val}} E \stackrel{\text{def}}{=} \underline{1} \leq E \wedge W \lesssim_{\text{list } A}^{\text{val}} E - \underline{1}
 \end{array}$$

Fig. 16. Extensions for handling call-by-value lists

```

rec sort(xs) =
case xs of nil ↦ nil
        | cons(y, ys) ↦ case ys of nil ↦ cons(y, nil)
                           | cons(z, zs) ↦ let q = divide(cons(y, ys)) in
                                           merge(sort(π1 q), sort(π2 q))

||sort|| = ⟨0, fix sort. λxs. if xs then ⟨0, 0⟩
           else if xs - 1 then ⟨0, 1⟩ else F(xs - 1)⟩
F(xs') = (2 ⊞ (divide(xs' + 1))c) +c E((divide(xs' + 1))p)
E(q) = (5 ⊞ (sort(π1 q))c ⊞ (sort(π2 q))c) +c merge ⟨(sort(π1 q))p, (sort(π2 q))p⟩

```

Fig. 17. The merge sort function in CBV PCF and its extracted recurrence

such as one that records both the size of the tree along with the maximum label. These different abstractions in turn allow for more or less detailed cost analyses of the original programs.

10.1 Merge Sort Example

We end with an example, viz. recurrence extraction for merge sort in CBV PCF. We assume the existence of two functions $\text{divide} : \text{list } A \rightarrow \text{list } A \times \text{list } A$ and $\text{merge} : \text{list } A \times \text{list } A \rightarrow \text{list } A$, which perform the usual tasks of splitting a list into two (nearly) equal-sized pieces and merging two sorted lists into a single sorted list. The translation into CBPV is extremely verbose, so we directly give the extracted recurrence in Figure 17. Corresponding to divide and merge are CBPV programs $\text{divide} : U(\text{list } A^* \rightarrow F(\text{list } A^* \times \text{list } A^*))$ and $\text{merge} : U(\text{list } A^* \times \text{list } A^* \rightarrow F(\text{list } A^*))$, and thus PCF_c recurrences $\text{divide} : \text{nat} \rightarrow \mathbb{C} \times \text{nat} \times \text{nat}$ and $\text{merge} : \text{nat} \times \text{nat} \rightarrow \mathbb{C} \times \text{nat}$. The recurrence for divide expresses the cost of the CBV PCF divide and the lengths of the two returned lists in terms of the length of its input. The recurrence extraction is, as with exponentiation, a tedious unwinding of definitions which uses a few additional size order axioms that are valid in the standard semantics of §8.

We can then analyze the semantic recurrence—i.e. the denotation of $\|\text{sort}\|$ —and the two projections for potential and cost. For this discussion, we will write syntactic expressions but manipulate them as though they are the corresponding denotations. If $S(n) = \pi_2(\|\text{sort}\|_p n)$ and $T(n) = \pi_1(\|\text{sort}\|_c n)$, then the Bounding Theorem tells us that $S(n)$ and $T(n)$ are upper bounds on the length and cost of $\text{sort}(xs)$ respectively, when xs has length n . We will assume that

$$\text{divide } k = \langle ck, \lceil k/2 \rceil, \lfloor k/2 \rfloor \rangle \quad \text{merge } \langle k, \ell \rangle = \langle d(k + \ell), k + \ell \rangle$$

for some fixed constants c and d . This is one way to formalize the assumption that the complexity of sort does not depend on the potentials of the actual list elements: while merge depends on a comparison function, we are asserting that its cost and the length of the result depend only on the lengths of its arguments. Of course, this is exactly what we typically do for an informal analysis of merge sort on lists of constant-size elements (e.g. machine words), or on lists of arbitrary data under the assumption that the comparison function takes constant time. Moving the abstraction of lists to numbers from the recurrence extraction to the denotational semantics would permit an analysis that takes into account the complexity of the comparison function.

We can now write out a more ‘traditional’ recurrence for S :

$$\begin{aligned} S(0) &= 0 & S(1) &= 1 & S(n) &= (E((\text{divide}(n))_p))_p = (E(\langle \lceil n/2 \rceil, \lfloor n/2 \rfloor \rangle))_p \\ & & & & &= (\text{merge } \langle S(\lceil n/2 \rceil), S(\lfloor n/2 \rfloor) \rangle)_p \\ & & & & &= S(\lceil n/2 \rceil) + S(\lfloor n/2 \rfloor) \end{aligned}$$

The standard proof tells us that $S(n) = n$. Next we write out a more traditional recurrence for $T(n)$ when n is a power of 2, which mimics the usual approach for analyzing such algorithms:

$$\begin{aligned} T(1) &= 0 & T(n) &= (2 + (\text{divide}(n))_c) + (E((\text{divide}(n))_p))_c \\ & & &= (2 + cn) + (E(\langle n/2, n/2 \rangle))_c \\ & & &= (2 + cn) + (5 + T(n/2) + T(n/2)) + (\text{merge } \langle S(n/2), S(n/2) \rangle)_c \\ & & &= 7 + cn + 2T(n/2) + (\text{merge } \langle n/2, n/2 \rangle)_c \\ & & &= 7 + (c + d)n + 2T(n/2). \end{aligned}$$

This is precisely the recurrence we expect. To proceed further with this example, we could define appropriate `divide` and `merge` functions and prove that the extracted recurrences satisfy the aforementioned assumptions.

11 RELATED WORK

There is a long history of techniques for extracting cost information from programs, probably starting with Wegbreit [1975], who computes closed bounds for Lisp programs. Much of the work has been done for imperative languages, as exemplified by the COSTA project for Java bytecode analysis [Albert et al. 2012, 2013] and SACO for parallel cost [Albert et al. 2018]. We cannot hope to review the entire field here, so concentrate on recent work on cost analysis for functional programs.

The idea that the potential of a function is itself a function is taken from Danner et al. [2015], who in turn adapt it from Danner and Royer [2007]. Danielsson [2008] focuses on amortized cost of lazy programs and makes use of what is essentially the writer monad $\mathbb{C} \times -$, though he requires the user to explicitly annotate programs with ticks (charges). The Resource Aware ML (RAML) project has achieved great results in automating the cost analysis of higher-order functional programs using automatic amortized resource analysis (AARA). This approach reduces the establishment of cost bounds on a program to type inference in a resource-aware type system, and that in turn is reduced to a linear programming problem. It has been used to verify the cost of significant portions of the OCaml libraries [Hoffmann et al. 2017] and for analyzing space usage in the presence of garbage collection [Niu and Hoffmann 2018], to name just a couple of recent achievements. However, the version of RAML that is current as of the time of this writing is restricted to deriving polynomial bounds, and hence it derives a quadratic bound for merge-sort; we simply derive a recurrence, and if the user of the system can prove a tighter bound on it than quadratic (e.g., by using the Master Theorem), more power to her. Another type-based approach is described by Avanzini and Dal Lago [2017], where a form of size types with index polymorphism is used. As with AARA, cost analysis comes down to type inference. This work also makes use of ticks (charges); an interesting aspect of it is that the clock itself becomes part of the program to which type inference is applied, and so cost comes down to an analysis of size. Somewhat earlier work takes a program transformation approach, defunctionalizing higher-order programs in a cost-preserving way, and then applying first-order analysis techniques to the result [Avanzini et al. 2015].

The accomplishments of these approaches in terms of automatically computing cost bounds is impressive, and not something we claim our approach as described here does. It is possible that after deploying our recurrence extraction technique, recurrence solvers such as OCRS [Kincaid et al. 2017] could be used, similarly to how RAML uses an off-the-shelf LP solver (of course,

since our recurrences are higher-order functions, this would presumably require some form of defunctionalization, but perhaps Benzinger’s [2004] work would be applicable). The main contrast between our work and all of the above is that our work puts the standard informal approaches to cost analysis via recurrence extraction on firm mathematical footing, while the automatic techniques use different methods than what we teach students in introductory classes or what programmers do in their heads.

12 CONCLUSIONS AND FUTURE WORK

We have improved upon the results of Danner et al. [2015] by giving a uniform recurrence extraction method for a variety of source languages. Instead of formulating extraction directly for each source language, we embed them in an intermediate language—namely CBPV—and perform recurrence extraction for that language instead. Our method uniformly handles general recursion irrespective of the evaluation strategy of the source language. We have shown this strategy in action by showing how to extract recurrences from both call-by-value and call-by-name PCF programs.

The natural next step would be to increase the expressiveness of the source language types. For example, it would be very interesting to examine whether our strategy can be extended to cover *recursive types* [Pierce 2002, §20]. This would lead us to an extraction function for call-by-value and call-by-name versions of Plotkin’s *Fixed Point Calculus (FPC)* [Plotkin 1985], [Gunter 1992, §7.4], thereby enabling the consideration of a large number of very expressive programming languages. On the other hand, such a programme would present significant technical challenges: amongst other things, it would require the solution of sized domain equations.

In either case, the combination of inductive or recursive data types and recursion would also aid us in obtaining a handle on the complexity of *coinductive data*. For example, in a call-by-value setting with inductive types and recursion we can express *streams* of natural numbers by the type $\mu X.1 \rightarrow (1 + \text{nat} \times X)$. General recursion then allows us to define non-trivial streams and functions that compute on them. Recurrence extraction in this setting would guide us towards various notions of *stream complexity*. Moreover, this would be a special case of *functional reactive programming (FRP)* [Elliott and Hudak 1997], and extending recurrence extraction to FRP would naturally lead us to interesting notions of complexity for functional programs that intentionally do not terminate.

In another direction, deterministic programs with interesting average-case behavior, such as those implementing Quicksort, typically use general recursion. It should be the case that the notions of probabilistic recurrences [Karp 1994] used to informally analyze such algorithms can be formalized in terms of appropriate semantic models in which the size of an argument is interpreted as an appropriate random variable. Finally, formalizing amortized analysis is another direction of significant interest. The current approach analyzes composition by composing worst-case bounds, which do not always yield a tight result. The goal would not be the automated analysis of AARA, but a formalization of informal techniques such as Tarjan’s [1985] banker’s and physicist’s methods.

ACKNOWLEDGMENTS

This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-16-1-0292. Any opinions, finding, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force. Additionally, this material is based upon work supported by the National Science Foundation under Grant Number 1618203.

REFERENCES

Samson Abramsky. 1990. The Lazy Lambda Calculus. In *Research Topics in Functional Programming*. Addison Wesley, 65–117. <https://www.cs.ox.ac.uk/files/293/lazy.pdf>

- Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. 1996. Full Abstraction for PCF. *Information and Computation* 163 (1996), 409–470.
- Samson Abramsky and Achim Jung. 1994. Domain Theory. In *Handbook of Logic in Computer Science*, Samson Abramsky, Dov M. Gabbay, and Thomas S. E. Maibaum (Eds.). Vol. 3. Oxford University Press, 1–168. <https://www.cs.bham.ac.uk/~axj/pub/papers/handy1.pdf>
- Elvira Albert, Puri Arenas, Samir Genaim, German Puebla, and Damiano Zanardini. 2012. Cost analysis of object-oriented bytecode programs. *Theoretical Computer Science* 413, 1 (2012), 142–159. <https://doi.org/10.1016/j.tcs.2011.07.009>
- Elvira Albert, Jesús Correas, Einar Broch Johnsen, Ka I. Pun, and Guillermo Román-Díez. 2018. Parallel Cost Analysis. *ACM Transactions on Computational Logic* 19, 4, Article 31 (Nov. 2018), 37 pages. <https://doi.org/10.1145/3274278>
- Elvira Albert, Samir Genaim, and Abu Naser Masud. 2013. On the Inference of Resource Usage Upper and Lower Bounds. *ACM Transactions on Computational Logic* 14, 3 (2013), 22:1–22:35. <https://doi.org/10.1145/2499937.2499943>
- Martin Avanzini and Ugo Dal Lago. 2017. Automating Sized-type Inference for Complexity Analysis. *Proceedings of the ACM on Programming Languages* 1, ICFP, Article 43 (2017), 29 pages. <https://doi.org/10.1145/3110287>
- Martin Avanzini, Ugo Dal Lago, and Georg Moser. 2015. Analyzing the complexity of functional programs: higher-order meets first-order. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming, ICFP 2015*, Kathleen Fisher and John Reppy (Eds.). 152–164. <https://doi.org/10.1145/2784731.2784753>
- Ralph Benzing. 2004. Automated higher-order complexity analysis. *Theoretical Computer Science* 318, 1-2 (2004), 79–103. <https://doi.org/10.1016/j.tcs.2003.10.022>
- Richard Bird. 2014. *Thinking Functionally with Haskell*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9781316092415>
- Guy Blelloch and John Greiner. 1995. Parallelism in Sequential Functional Languages. In *Proceedings of the Seventh International Conference on Functional Programming Languages and Computer Architecture (FPCA '95)*. ACM, New York, NY, USA, 226–237. <https://doi.org/10.1145/224164.224210>
- Nils Anders Danielsson. 2008. Lightweight semiformal time complexity analysis for purely functional data structures. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, George Necula and Philip Wadler (Eds.). 133–144. <https://doi.org/10.1145/1328438.1328457>
- Norman Danner, Daniel R. Licata, and Ramyaa Ramyaa. 2015. Denotational cost semantics for functional languages with inductive types. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming - ICFP 2015*. ACM Press, New York, New York, USA, 140–151. <https://doi.org/10.1145/2784731.2784749>
- Norman Danner, Jennifer Paykin, and James S. Royer. 2013. A static cost analysis for a higher-order language. In *Proceedings of the 7th workshop on Programming languages meets program verification - PLPV '13*. ACM Press, New York, New York, USA. <https://doi.org/10.1145/2428116.2428123>
- Norman Danner and James S. Royer. 2007. Adventures in time and space. *Logical Methods in Computer Science* 3, 9 (2007), 1–53. [https://doi.org/10.2168/LMCS-3\(1:9\)2007](https://doi.org/10.2168/LMCS-3(1:9)2007)
- Conal Elliott and Paul Hudak. 1997. Functional reactive animation. In *Proceedings of the second ACM SIGPLAN International Conference on Functional Programming - ICFP '97*. ACM Press, New York, New York, USA, 263–273. <https://doi.org/10.1145/258948.258973>
- Martín Hötzel Escardó and Weng Kin Ho. 2009. Operational domain theory and topology of sequential programming languages. *Information and Computation* 207, 3 (2009), 411–437. <https://doi.org/10.1016/j.ic.2008.12.003>
- Marcelo P. Fiore. 1996. *Axiomatic Domain Theory in Categories of Partial Maps*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9780511526565>
- Carl A. Gunter. 1992. *Semantics of Programming Languages: Structures and Techniques*. The MIT Press.
- Jan Hoffmann, Ankush Das, and Shu-Chun Weng. 2017. Towards Automatic Resource Bound Analysis for OCaml. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL 2017)*. ACM, New York, NY, USA, 359–373. <https://doi.org/10.1145/3009837.3009842>
- Richard M. Karp. 1994. Probabilistic recurrence relations. *Journal of the Association for Computing Machinery* 41, 6 (1994), 1136–1150. <https://doi.org/10.1145/195613.195632>
- Theodore Seok Kim. 2016. *Cost Semantics for Plotkin's PCF*. Master's Thesis. Wesleyan University. https://wesscholar.wesleyan.edu/etd_mas_theses/130/
- Zachary Kincaid, John Cyphert, Jason Breck, and Thomas Reps. 2017. Non-linear Reasoning for Invariant Synthesis. *Proceedings of the ACM on Programming Languages* 2, POPL, Article 54 (Dec. 2017), 33 pages. <https://doi.org/10.1145/3158142>
- Paul Blain Levy. 2003. *Call-by-Push-Value: A Functional-Imperative Synthesis*. Springer. <https://doi.org/10.1007/978-94-007-0954-6>
- Paul Blain Levy. 2006. Call-by-push-value: Decomposing call-by-value and call-by-name. *Higher-Order and Symbolic Computation* 19 (2006), 377–414. <https://doi.org/10.1007/s10990-006-0480-6>

- Robin Milner. 1977. Fully abstract models of typed λ -calculi. *Theoretical Computer Science* 4, 1 (1977), 1–22. [https://doi.org/10.1016/0304-3975\(77\)90053-6](https://doi.org/10.1016/0304-3975(77)90053-6)
- John C. Mitchell. 1996. *Foundations for programming languages*. The MIT Press.
- Yue Niu and Jan Hoffmann. 2018. Automatic Space Bound Analysis for Functional Programs with Garbage Collection. In *LPAR-22. 22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning (EPIc Series in Computing)*, Gilles Barthe, Geoff Sutcliffe, and Margus Veanes (Eds.), Vol. 57. EasyChair, 543–563. <https://doi.org/10.29007/xkwx>
- Benjamin C. Pierce. 2002. *Types and Programming Languages*. The MIT Press.
- A. M. Pitts. 1994. *Some notes on inductive and co-inductive techniques in the semantics of functional programs*. Technical Report. BRICS. <https://www.brics.dk/NS/94/5/BRICS-NS-94-5.pdf>
- A. M. Pitts. 1997. Operationally-Based Theories of Program Equivalence. In *Semantics and Logics of Computation*, A. M. Pitts and P. Dybjer (Eds.). Cambridge University Press, 241–298. <https://www.cl.cam.ac.uk/~amp12/papers/oepbtp/oepbtp.pdf>
- Gordon Plotkin and John Power. 2001. Adequacy for Algebraic Effects. In *Foundations of Software Science and Computation Structures. FoSSaCS 2001*, Furio Honsell and M. Miculan (Eds.). Lecture Notes in Computer Science, Vol. 2030. Springer, Berlin, Heidelberg. https://doi.org/10.1007/3-540-45315-6_1
- Gordon D. Plotkin. 1977. LCF considered as a programming language. *Theoretical Computer Science* 5, 3 (1977), 223–255. [https://doi.org/10.1016/0304-3975\(77\)90044-5](https://doi.org/10.1016/0304-3975(77)90044-5)
- G. D. Plotkin. 1985. Lectures on predomains and partial functions. Notes for a course given at CSLI, Stanford University.
- Jon G. Riecke. 1993. Fully abstract translations between functional languages. *Mathematical Structures in Computer Science* 3 (1993), 387–415. <https://doi.org/10.1017/S0960129500000293>
- V. Stoltenberg-Hansen, I. Lindstrom, and E. R. Griffor. 1994. *Mathematical Theory of Domains*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9781139166386>
- Thomas Streicher. 2006. *Domain-theoretic Foundations of Functional Programming*. World Scientific.
- Robert Endre Tarjan. 1985. Amortized computational complexity. *SIAM J. Algebraic Discrete Methods* 6, 2 (1985), 306–318. <https://doi.org/10.1137/0606031>
- Ben Wegbreit. 1975. Mechanical program analysis. *Communications of the Association for Computing Machinery* 18, 9 (1975), 528–539. <https://doi.org/10.1145/361002.361016>