

Denotational Cost Semantics for Functional Languages with Inductive Types

Norman Danner*

Wesleyan University, USA
ndanner@wesleyan.edu

Daniel R. Licata[†]

Wesleyan University, USA
dlicata@wesleyan.edu

Ramyaa Ramyaa

Wesleyan University, USA
ramyaa@wesleyan.edu

Abstract

A central method for analyzing the asymptotic complexity of a functional program is to extract and then solve a recurrence that expresses evaluation cost in terms of input size. The relevant notion of input size is often specific to a datatype, with measures including the length of a list, the maximum element in a list, and the height of a tree. In this work, we give a formal account of the extraction of cost and size recurrences from higher-order functional programs over inductive datatypes. Our approach allows a wide range of programmer-specified notions of size, and ensures that the extracted recurrences correctly predict evaluation cost. To extract a recurrence from a program, we first make costs explicit by applying a monadic translation from the source language to a complexity language, and then abstract datatype values as sizes. Size abstraction can be done semantically, working in models of the complexity language, or syntactically, by adding rules to a preorder judgement. We give several different models of the complexity language, which support different notions of size. Additionally, we prove by a logical relations argument that recurrences extracted by this process are upper bounds for evaluation cost; the proof is entirely syntactic and therefore applies to all of the models we consider.

Categories and Subject Descriptors F.3.1 [Logics and meanings of programs]: Specifying and verifying and reasoning about programs; F.3.2 [Logics and meanings of programs]: Semantics of programming languages

General Terms Verification.

Keywords Semi-automatic complexity analysis.

* This material is based upon work supported by the National Science Foundation under grant no. 1318864.

[†] This material is based on research sponsored by The United States Air Force Research Laboratory under agreement number FA9550-15-1-0053. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the United States Air Force Research Laboratory, the U.S. Government or Carnegie Mellon University.

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

ICFP'15, August 31 – September 2, 2015, Vancouver, BC, Canada.
Copyright © 2015 ACM 978-1-4503-3669-7/15/08...\$15.00.
<http://dx.doi.org/10.1145/nnnnnnn.nnnnnn>

1. Introduction

The typical method for analyzing the asymptotic complexity of a functional program is to extract a recurrence that relates the function's running time to the size of the function's input, and then solve the recurrence to obtain a closed form and big- O bound. Automated complexity analysis (see the related work in Section 7) provides helpful information to programmers, and could be particularly useful for giving feedback to students. In a setting with higher-order functions and programmer-defined datatypes, automating the extract-and-solve method requires a generalization of the standard theory of recurrences. This generalization must include a notion of recurrence for higher-order functions such as `map` and `fold`, as well as a general theory of what constitutes “the size of the input” for programmer-defined datatypes.

One notion of recurrence for higher-order functions was developed in previous work by Danner and Royer (2009) and Danner et al. (2013). Because the output of one function is the input to another, it is necessary to extract from a function not only a recurrence for the running time, but also a recurrence for the size of the output. These can be packaged together as a single recurrence that, given the size of the input, produces a pair consisting of the running time (called the *cost*) and the size of the output (called the *potential*). Whereas the former is the cost of executing the program to a value, the latter determines the cost of using that value. This generalizes naturally to higher-order functions: a recurrence for a higher-order function is itself a higher-order function, which expresses the cost and potential of the result in terms of a given recurrence for the cost and potential of the argument function. The process of extracting recurrences can thus be seen as a denotational semantics of the program, where a function is interpreted as a function from input potential to cost and output potential.

Building on this work, we give a formal account of the extraction of recurrences from higher-order functional programs over inductive datatypes, focusing how to soundly allow programmer-specified sizes of datatypes. We show that under some mild conditions on sizes, the cost predicted by an extracted recurrence is in fact an upper bound on the number of steps the program takes to evaluate. The size of a value can be taken to be (essentially) the value itself, in which case one gets exact bounds but must reason about all the details of program evaluation, or the size of a value can forget information (e.g. abstracting a list as its length), in which case one gets weaker bounds with more traditional reasoning.

We start from a call-by-value source language, defined in Section 2, with strictly positive inductive datatype definitions (which include lists and finitely and infinitely branching trees) Datatypes are used via case-analysis and structural recursion (so the language is terminating), but unlike in Danner et al. (2013), recursive calls are only evaluated if necessary—for example, recurring on one branch of a tree has different cost than recurring on both branches. The

cost of a program is defined by an operational cost semantics, an evaluation relation annotated with costs. For simplicity, the cost semantics measures only the number of function applications and recursive calls made during evaluation, but our approach to extracting recurrences generalizes to other cost models.

We extract a recurrence from such a program in two steps. First, in Section 3, we make the cost of evaluating a program explicit, by translating a source program e to a program $\|e\|$ in a complexity language. The complexity language has an additional type \mathbf{C} for costs, and the translation to the complexity language is a call-by-value monadic translation into the writer monad $\mathbf{C} \times -$ (Moggi 1991; Wadler 1992). The translated program $\|e\|$ returns an additional result, which is the cost of running the original program e .

Second, we abstract values to sizes; we study both semantic and syntactic approaches. In Section 4, we give a size-based semantics of the complexity language, which relies on programmer-specified size functions mapping each datatype to the natural numbers (or some other preorder). Typical size functions include the length of a list and the size or depth of a tree. The semantics satisfies a *bounding theorem* (Theorem 7), which implies that the denotational cost given by composing the source-to-complexity translation with the size-based semantics is in fact an upper bound on the operational cost. We show some examples that the recurrence for cost extracted by this process is the expected one; later we also show that the results in Danner et al. (2013) carry over.

Alternatively, the abstraction of values to sizes can be done syntactically in the complexity language, by imposing a preorder structure on the values of the datatype themselves. For example, rather than mapping lists to numbers representing their lengths, we can order the list values by rules including $xs \leq (x::xs)$ and $(x::xs) \leq (y::xs)$. The second rule says that the elements of the list are irrelevant, quotienting the lists down to natural numbers, and the first generates the usual order on natural numbers. Formally, we equip the complexity language with a judgement $E \leq E'$ that can be used to make such abstractions. In Section 5, we identify properties of this judgement that are sufficient to prove a syntactic bounding theorem (Theorem 12), which states that the operational cost is bounded by the cost component of the complexity translation. The key technical notion is a logical relation between the source and complexity languages that extends the bounding relation of Danner et al. (2013) to inductive types. This proof gives a bounding theorem for any model of the complexity language that validates the rules for \leq . In Section 6, we show that these rules are valid in the size-based semantics of Section 4 (thereby proving Theorem 7), and we discuss several other models of the complexity language.

This gives a formal account of what it means to extract a recurrence from higher-order programs on inductive data. We leave an investigation of what it means to solve these higher-order recurrences to future work. Danner et al. (2015) is a full version of this paper.

2. Source Language with Inductive Data Types

The source language is a simply-typed λ -calculus with product types, function types, suspensions, and strictly positive inductive datatypes. Its syntax, typing, and operational semantics are given in Figure 1. We bundle sums and inductive types together as datatypes, rather than using separate $+$ and μ types, because below we do not want to consider sizes for the sum part separately. We assume a top-level signature ψ consisting of datatype declarations of the form

$$\text{datatype } \delta = C_0^\delta \text{ of } \phi_{C_0}[\delta] \mid \dots \mid C_{n-1}^\delta \text{ of } \phi_{C_{n-1}}[\delta]$$

Each constructor’s argument type is specified by a strictly positive functor ϕ . These include the identity functor (t), representing a recursive occurrence of the datatype; constant functors (τ), representing a non-recursive argument; product functors ($\phi_1 \times \phi_2$), representing a

pair of arguments; and constant exponentials ($\tau \rightarrow \phi$), representing an argument of function type. For example for τ `list`, the argument type for `Nil` is `unit` (constant functor), and the argument type for `Cons` is $\tau \times t$ (product of constant and recursive arguments). We write $\phi[\tau] = \phi[\tau/t]$ for substitution for the single free type variable t in ϕ . We sometimes abbreviate further by dropping the type superscripts and writing `datatype` $\delta = C$ of ϕ_C and by writing C rather than C_i to refer to one of the constructors of the declaration. In the signature, each ϕ_C in each `datatype` declaration must refer only to datatypes that are declared earlier in the sequence, to avoid introducing general recursive datatypes (see the full paper for the formal definition). We write $C : (\phi \rightarrow \delta) \in \psi$ to mean that the signature ψ contains a datatype declaration of the form `datatype` $\delta = \dots \mid C$ of $\phi[\delta] \mid \dots$. We generally elide the signature from typing, but sometimes write $\gamma \vdash_\psi e : \tau$ to include it. The elimination rule for a datatype δ is structural recursion, $\text{rec}^\delta(e, \overline{C} \mapsto x.e_C)$. When $\phi_C = \text{unit}$, we assume $x \notin \text{fv}(e_C)$ and write e_C instead of $x.e_C$.

Evaluation is call-by-value and products and datatypes are strict. However, unfolding datatype recursors requires substituting expressions (the recursor applied to the components of the value) for the variables standing for the recursive calls—running the recursive call first and substituting its value would require a function to make all possible recursive calls. We handle this using suspensions: when computing a τ by recursion, the result of a recursive call is given the type `susp` τ . The values of type `susp` τ are `delay`(e) where e is an expression of type τ ; the elimination form `force` forces evaluation. When defining a recursive computation of result type τ , the branch for a constructor C has access to a variable of type $\phi_C[\delta \times \text{susp } \tau]$, which gives access both to the “predecessor” values of type δ and to the recursive results. This supports both case-analysis and structural recursion, and recursive calls are only computed if they are used.

For any strictly positive functor ϕ , the `map` ^{ϕ} expression witnesses functoriality, essentially lifting a function $\tau_0 \rightarrow \tau_1$ to a function $\phi[\tau_0] \rightarrow \phi[\tau_1]$. It is used in the operational semantics for the recursor to insert recursive calls at the right places in ϕ (Harper (2013) provides an exposition). We will only need to lift maps $x : \tau_0.v : \tau_1$ whose bodies are syntactic values (or variables), and apply them to syntactic values (or variables), and we restrict `map` to this special case to simplify its cost semantics.

The cost semantics in Figure 1 defines the relation $e \downarrow^n v$, which means that the expression e evaluates to the value v in n steps. Our cost model charges only for the number of function applications and recursive calls made by datatype recursors. This prevents constant-time overheads from the encoding of datatypes using product and suspension types from showing up in the extracted recurrences. It is simple to adapt the denotational cost semantics below to other operational cost semantics, such as one that charges for these steps, or assigns different costs to different constructs.

Substitutions are defined as usual:

DEFINITION 1. We write θ for substitutions $v_1/x_1, \dots, v_n/x_n$, and $\theta : \gamma$ to mean that $\text{Dom } \theta \subseteq \text{Dom } \gamma$ and $\theta \vdash \theta(x) : \gamma(x)$ for all $x \in \text{Dom } \theta$. We define the application of a substitution θ to an expression e as usual and denote it $e[\theta]$.

LEMMA 1. If $x \notin \text{Dom } \theta$, then $e[\theta, x/x][e_1/x] = e[\theta, e_1/x]$.

For source cost expressions n , we write $n \leq n'$ for the order given by interpreting these cost expressions as natural numbers (i.e. the free precongruence generated by the monoid equations for $(+, 0)$ and $0 \leq 1$). We have the following syntactic properties of evaluation:

LEMMA 2 (Value Evaluation).

- If $v \downarrow^n v'$ then $n \leq 0$ and $v = v'$.

Types:

$$\begin{aligned} \tau &::= \text{unit} \mid \tau \times \tau \mid \tau \rightarrow \tau \mid \text{susp } \tau \mid \delta \\ \phi &::= t \mid \tau \mid \phi \times \phi \mid \tau \rightarrow \phi \\ \text{datatype } \delta &= C_0^\delta \text{ of } \phi_{C_0}[\delta] \mid \dots \mid C_{n-1}^\delta \text{ of } \phi_{C_{n-1}}[\delta] \end{aligned}$$

Expressions:

$$\begin{aligned} v &::= x \mid \langle \rangle \mid \langle v, v \rangle \mid \lambda x.e \mid \text{delay}(e) \mid C v \\ e &::= x \mid \langle \rangle \mid \langle e, e \rangle \mid \text{split}(e, x.x.e) \mid \lambda x.e \mid e e \\ &\quad \mid \text{delay}(e) \mid \text{force}(e) \\ &\quad \mid C^\delta e \mid \text{rec}^\delta(e, \overline{C} \mapsto x.e_C) \\ &\quad \mid \text{map}^\phi(x.v, v) \mid \text{let}(e, x.e) \\ n &::= 0 \mid 1 \mid n + n \end{aligned}$$

Operational semantics: $e \downarrow^n v$.

$$\frac{}{\text{delay}(e) \downarrow^0 \text{delay}(e)} \quad \frac{e \downarrow^{n_0} \text{delay}(e_0) \quad e_0 \downarrow^{n_1} v}{\text{force}(e) \downarrow^{n_0+n_1} v} \quad \frac{e \downarrow^n v}{C e \downarrow^n C v}$$

$$\frac{e \downarrow^{n_0} C v_0 \quad \text{map}^{\phi_C}(y. \langle y, \text{delay}(\text{rec}(y, \overline{C} \mapsto x.e_C)) \rangle, v_0) \downarrow^{n_1} v_1 \quad e_C[v_1/x] \downarrow^{n_2} v}{\text{rec}(e, \overline{C} \mapsto x.e_C) \downarrow^{1+n_0+n_1+n_2} v}$$

Typing: $\gamma \vdash e : \tau$.

$$\frac{\gamma \vdash e : \tau}{\gamma \vdash \text{delay}(e) : \text{susp } \tau} \quad \frac{\gamma \vdash e : \text{susp } \tau}{\gamma \vdash \text{force}(e) : \tau}$$

$$\frac{\gamma \vdash e : \phi_C[\delta]}{\gamma \vdash C^\delta e : \delta}$$

$$\frac{\gamma \vdash e : \delta \quad \forall C (\gamma, x : \phi_C[\delta \times \text{susp } \tau] \vdash e_C : \tau)}{\gamma \vdash \text{rec}^\delta(e, \overline{C} \mapsto x.e_C) : \tau}$$

$$\frac{\gamma, x : \tau_0 \vdash v_1 : \tau_1 \quad \gamma \vdash v_0 : \phi[\tau_0]}{\gamma \vdash \text{map}^\phi(x.v_1, v_0) : \phi[\tau_1]}$$

$$\frac{}{\text{map}^t(x.v, v_0) \downarrow^0 v[v_0/x]} \quad \frac{}{\text{map}^\tau(x.v, v_0) \downarrow^0 v_0} \quad (t \text{ not free in } \tau)$$

$$\frac{\text{map}^{\phi_0}(x.v, v_0) \downarrow^{n_0} v'_0 \quad \text{map}^{\phi_1}(x.v, v_1) \downarrow^{n_1} v'_1}{\text{map}^{\phi_0 \times \phi_1}(x.v, \langle v_0, v_1 \rangle) \downarrow^{n_0+n_1} \langle v'_0, v'_1 \rangle}$$

$$\frac{}{\text{map}^{\tau \rightarrow \phi}(x.v, \lambda y.e) \downarrow^0 \lambda y. \text{let}(e, z. \text{map}^\phi(x.v, z))}$$

Figure 1: Source language syntax and typing and operational semantics. Standard typing rules for variables, product types, function types, and `let` are omitted. In omitted operational rules, the costs are the sum of the costs of the subevaluations, except for e_0 e_1 , which adds 1.

- For all $v, v \downarrow^0 v$.

LEMMA 3 (Totality of `map`). If $\gamma \vdash \text{map}^\phi(x.v_1, v_0) : \phi[\tau_1]$ then $\text{map}^\phi(x.v_1, v_0) \downarrow^0 v$ for some v .

3. Making Costs Explicit

3.1 The Complexity Language

The complexity language will serve as a monadic metalanguage (Moggi 1991) in which we make evaluation cost explicit. The syntax and typing are given in Figure 2. The preorder judgement defined in Section 5 will play a role analogous to an operational or equational semantics for the complexity language.

Because we are not concerned with the evaluation steps of the complexity language itself, we remove features of the source language that were used to control evaluation costs. Product types are eliminated by projections, rather than `split`. We allow substitution of arbitrary expressions for variables, which is used in recursors for datatypes. Consequently, suspensions are not necessary. We treat $\text{map}^\Phi(x.E, E_1)$ as an admissible rule (macro), defined by induction on Φ :

$$\frac{\Gamma, x : T_0 \vdash E_1 : T_1 \quad \Gamma \vdash E_0 : \Phi[T_0]}{\Gamma \vdash \text{map}^\Phi(x.E_1, E_0) : \Phi[T_1]}$$

$$\text{map}^t(x.E, E_0) := E[E_0/x]$$

$$\text{map}^T(x.E, E_0) := E_0$$

$$\text{map}^{\Phi_0 \times \Phi_1}(x.E, E_0) := \langle \text{map}^{\Phi_0}(x.E, \pi_0 E_0), \text{map}^{\Phi_1}(x.E, \pi_1 E_0) \rangle$$

$$\text{map}^{T \rightarrow \Phi}(x.E, E_1) := \lambda y. \text{map}^\Phi(x.E, E_1 y)$$

The type \mathbf{C} represents some domain of costs. The term constructors for \mathbf{C} say only that it is a monoid $(+, 0)$ with a value 1 representing a single step. Costs can be interpreted in a variety of

ways—e.g. as natural numbers and as natural numbers with infinity (Section 4).

Substitutions Θ in the complexity language are defined as usual, and satisfy standard composition properties:

LEMMA 4.

- If x does not occur in Θ , then $E[\Theta, x/x][E_1/x] = E[\Theta, E_1/x]$.
- If x_1, x_2 do not occur in Θ , then $E[E_1/x_1][E_2/x_2][\Theta] = E[\Theta, E_1[\Theta]/x_1, E_2[\Theta]/x_2]$.

3.2 The Complexity Translation

Consider a higher-order function such as `map` on lists:

$$\text{listmap} = \lambda(f, xs). \text{rec}(xs, \text{Nil} \mapsto \text{Nil} \mid \text{Cons} \mapsto \langle y, \langle ys, r \rangle \rangle. \text{Cons}(f y, \text{force}(r)))$$

The cost of $\text{listmap}(f, xs)$ depends on the sizes of each element of xs , and the cost of evaluating f on elements of those sizes. However, since $\text{listmap}(f, xs)$ might itself be an argument to another function (e.g. another `listmap`), we also need to predict the sizes of the elements of $\text{listmap}(f, xs)$, which depends on the size of the output of f . Thus, to analyze `listmap`, we should be given a recurrence for the cost and size of $f(x)$ in terms of the size of x , and need to produce a recurrence that gives the cost and size of $\text{listmap}(f, xs)$ in terms of the size of xs . We call the size of the value of an expression that expression's *potential*, because the size of the value determines what future uses of that value will cost.

This discussion motivates a translation $\|\cdot\|$ from source language terms to complexity language terms so that if $e : \tau$, then $\|e\| : \mathbf{C} \times \langle \tau \rangle$. In the complexity language, we call an expression of type $\mathbf{C} \times \langle \tau \rangle$ a *complexity*, an expression of type \mathbf{C} a *cost*, and an expression

Types:

$$\begin{aligned} T & ::= \mathbf{C} \mid \text{unit} \mid \Delta \mid T \times T \mid T \rightarrow T \\ \Phi & ::= t \mid T \mid \Phi \times \Phi \mid T \rightarrow \Phi \\ \text{datatype } \Delta & = C_0^\Delta \text{ of } \Phi_{C_0}[\Delta] \mid \dots \mid C_{n-1}^\Delta \text{ of } \Phi_{C_{n-1}}[\Delta] \end{aligned}$$

Expressions:

$$\begin{aligned} E & ::= x \mid 0 \mid 1 \mid E + E \mid \\ & \langle \rangle \mid \langle E, E \rangle \mid \pi_0 E \mid \pi_1 E \mid \lambda x. E \mid E E \\ & \mid C^\Delta E \mid \text{rec}^\Delta(E, \bar{C} \mapsto x. \bar{E}_C) \end{aligned}$$

Typing: $\Gamma \vdash E : T$.

$$\frac{\Gamma \vdash E : \Phi_C[\Delta]}{\Gamma \vdash C^\Delta E : \Delta}$$

$$\frac{\Gamma \vdash E : \Delta \quad \forall C (\Gamma, x : \Phi_C[\Delta \times T] \vdash E_C : T)}{\Gamma \vdash \text{rec}^\Delta(E, \bar{C} \mapsto x. \bar{E}_C) : T}$$

Figure 2: Complexity language types, expressions, and typing.

The typing rules are standard for unit, product, and arrow types. \mathbf{C} has a binary operation $+$ and elements 0 and 1.

of type $\langle\langle\tau\rangle\rangle$ a *potential*. We abbreviate $\mathbf{C} \times \langle\langle\tau\rangle\rangle$ by $\|\tau\|$. The first component of $\|e\|$ is the cost of evaluating e , and the second component of $\|e\|$ is the potential of e .

For example, `listmap` is a value, so its cost should be zero. On the other hand, its potential should describe what future uses of `listmap` will cost, in terms of the potentials of its arguments. For the type of `listmap` (uncurried), the above discussion suggests

$$\begin{aligned} \langle\langle\tau \rightarrow \sigma\rangle\rangle \times \langle\langle\tau \text{ list}\rangle\rangle & \rightarrow \langle\langle\sigma \text{ list}\rangle\rangle := \\ \langle\langle\tau\rangle\rangle \rightarrow \mathbf{C} \times \langle\langle\sigma\rangle\rangle & \times \langle\langle\tau \text{ list}\rangle\rangle \rightarrow \mathbf{C} \times \langle\langle\sigma \text{ list}\rangle\rangle \end{aligned}$$

For the argument function, we are provided a recurrence that maps τ -potentials to costs and σ -potentials. For the argument list, we are provided a τ `list`-potential. Using these, the potential of `listmap` must give the cost for doing the whole `map` and give a σ `list`-potential for the value. This illustrates how the potential of a higher-order function is itself a higher-order function.

As discussed above, we stage the extraction of a recurrence, and in the first phase, we do not abstract values as sizes (e.g. we do not replace a list by its length). Because of this, the complexity translation has a succinct description. For any monoid $(\mathbf{C}, +, 0)$, the writer monad (Wadler 1992) $\mathbf{C} \times -$ is a monad with

$$\begin{aligned} \text{return}(E) & := (0, E) \\ E_1 \gg E_2 & := (\pi_0(E_1) + \pi_0(E_2(\pi_1(E_1))), \pi_1(E_2(\pi_1(E_1)))) \end{aligned}$$

The monad laws follow from the monoid laws for \mathbf{C} . Thinking of \mathbf{C} as costs, these say that the cost of `return`(e) is zero, and that the cost of `bind` is the sum of the cost of E_1 and the cost of E_2 on the potential of E_1 . The complexity translation is then a call-by-value monadic translation from the source language into the writer monad in the complexity language, where source expressions that cost a step have the “effect” of incrementing the cost component, using the monad operation

$$\text{incr}(E : \mathbf{C}) : \mathbf{C} \times \text{unit} := (E, \langle \rangle)$$

We write this translation out explicitly in Figure 3. When E is a complexity, we write E_c and E_p for $\pi_0 E$ and $\pi_1 E$ respectively (for “cost” and “potential”). We will often need to “add cost” to a complexity; when E_1 is a cost and E_2 a complexity, we write $E_1 +_c E_2$ for the complexity $(E_1 + (E_2)_c, (E_2)_p)$ (in monadic notation, $\text{incr}(E_1) \gg E_2$). The type translation is extended pointwise to contexts, so $x : \tau \in \gamma$ iff $x : \langle\langle\tau\rangle\rangle \in \langle\langle\gamma\rangle\rangle$ —the translation is call-by-value, so variables range over potentials, not complexities. For example, $\|x\| = (0, x)$, where the x on the left is a source variable and the x on the right is a potential variable. Likewise we assume that for every datatype δ in the source signature, we have a corresponding datatype δ declared in the complexity language.

We note some basic facts about the translation: the type translation commutes with the application of a strictly positive functor, which is used to show that the translation preserves types.

LEMMA 5 (Compositionality).

- $\|\phi[\tau]\| = \|\phi\|[\langle\langle\tau\rangle\rangle]$
- $\langle\langle\phi[\tau]\rangle\rangle = \langle\langle\phi\rangle\rangle[\langle\langle\tau\rangle\rangle]$

THEOREM 6. If $\gamma \vdash_\psi e : \tau$, then $\|\gamma\| \vdash_{\|\psi\|} \|e\| : \|\tau\|$.

4. A Size-Based Complexity Semantics

In the above translation, the potential of a value has just as much information as that value itself. Next, we investigate how to abstract values to sizes, such as replacing a list by its length. In this section, we make this replacement by defining a size-based denotational semantics of the complexity language.

We need to be able to treat potentials of inductively-defined data in two different ways. On the one hand, potentials must reflect intuitions about sizes. To that end, we will insist that potentials be partial orders. On the other hand, to interpret `rec` expressions, we must be able to distinguish the datatype constructor that a potential represents. In other words, we need the potentials to also be (something like) inductive data types. We will have our cake and eat it too using an approach similar to the work on views (Wadler 1987). As hinted above, we interpret each datatype Δ in the complexity language as a partial order $\llbracket\Delta\rrbracket$. But we will also make use of the sum type $D^\Delta = \llbracket\Phi_{C_0}[\Delta]\rrbracket + \dots + \llbracket\Phi_{C_{n-1}}[\Delta]\rrbracket$ (representing the unfolding of the datatype) and a function $\text{size}_\Delta : D^\Delta \rightarrow \llbracket\Delta\rrbracket$ (which represents the size of a constructor, in terms of the size of the argument to the constructor). When $\Phi_{C_i} = t$ (i.e. the argument to the constructor is a single recursive occurrence of the datatype), $\text{size}(inj; x)$ is intended to represent an upper bound on the size of the values of the form Cv , where v is a value of size at most x . To define the semantics of $\text{rec}^\Delta(y, \bar{C} \mapsto x. \bar{E}_C)$, we consider all values $z \in D^\Delta$ such that $\text{size}_\Delta(z) \leq y$. We can distinguish between such values to (recursively) compute the possible values of the form $E_C[\dots/x]$, and then take a maximum over all such values.

For example, for the inductive definitions of `nat` and `list` (where the list elements have type `nat`), suppose we want to construe the size of a `list` to be the number of all `nat` and `list` constructors. We implement this in the complexity semantics as

$$\begin{aligned} \llbracket\text{nat}\rrbracket & = \mathbf{Z}^+ \\ D^{\text{nat}} & = \{\ast\} + \llbracket\text{nat}\rrbracket \\ \text{size}_{\text{nat}}(\ast) & = 1 \\ \text{size}_{\text{nat}}(m) & = 1 + m \\ \llbracket\text{list}\rrbracket & = \mathbf{Z}^+ \\ D^{\text{list}} & = \{\ast\} + (\llbracket\text{nat}\rrbracket \times \llbracket\text{list}\rrbracket) \\ \text{size}_{\text{list}}(\ast) & = 1 \\ \text{size}_{\text{list}}((m, n)) & = 1 + m + n \end{aligned}$$

where \mathbf{Z}^+ is the non-negative integers.

$$\begin{array}{lcl}
\|\tau\| & = & \mathbf{C} \times \langle\langle\tau\rangle\rangle \\
\langle\langle\mathbf{unit}\rangle\rangle & = & \mathbf{unit} \\
\langle\langle\sigma \times \tau\rangle\rangle & = & \langle\langle\sigma\rangle\rangle \times \langle\langle\tau\rangle\rangle \\
\langle\langle\sigma \rightarrow \tau\rangle\rangle & = & \langle\langle\sigma\rangle\rangle \rightarrow \langle\langle\tau\rangle\rangle \\
\langle\langle\mathbf{susp} \tau\rangle\rangle & = & \|\tau\| \\
\langle\langle\delta\rangle\rangle & = & \delta \\
\\
\|\phi\| & = & \mathbf{C} \times \langle\langle\phi\rangle\rangle \\
\langle\langle t \rangle\rangle & = & t \\
\langle\langle \tau \rangle\rangle & = & \langle\langle \tau \rangle\rangle \\
\langle\langle \phi_0 \times \phi_1 \rangle\rangle & = & \langle\langle \phi_0 \rangle\rangle \times \langle\langle \phi_1 \rangle\rangle \\
\langle\langle \tau \rightarrow \phi \rangle\rangle & = & \langle\langle \tau \rangle\rangle \rightarrow \langle\langle \phi \rangle\rangle
\end{array}$$

$\langle\langle\psi\rangle\rangle$ has, for each datatype δ in ψ
datatype $\delta = C_0^\delta$ of $\langle\langle\phi_{C_0}\rangle\rangle[\delta], \dots, C_{C_{n-1}}^\delta$ of $\langle\langle\phi_{n-1}\rangle\rangle[\delta]$

$$\begin{array}{lcl}
\|x\| & = & \langle 0, x \rangle \\
\|\langle \rangle\| & = & \langle 0, \langle \rangle \rangle \\
\|\langle e_0, e_1 \rangle\| & = & \langle \|e_0\|_c + \|e_1\|_c, \langle \|e_0\|_p, \|e_1\|_p \rangle \rangle \\
\|\mathbf{split}(e_0, x_0.x_1.e_1)\| & = & \|e_0\|_c + \|e_1\|_c \langle \pi_0 \|e_0\|_p / x_0, \pi_1 \|e_1\|_p / x_1 \rangle \\
\|\lambda x.e\| & = & \langle 0, \lambda x.e \rangle \\
\|e_0 e_1\| & = & (1 + (e_0)_c + (e_1)_c) +_c (e_0)_p (e_1)_p \\
\|\mathbf{delay}(e)\| & = & \langle 0, \|e\| \rangle \\
\|\mathbf{force}(e)\| & = & \|e\|_c +_c \|e\|_p \\
\|C_i^\delta e\| & = & \langle \|e\|_c, C_i^\delta \|e\|_p \rangle \\
\|\mathbf{rec}^\delta(e, \overline{C \mapsto x.e_C})\| & = & \|e\|_c +_c \mathbf{rec}^\delta(\|e\|_p, \overline{C \mapsto x.1 +_c \|e_C\|}) \\
\|\mathbf{map}^\phi(x.v_0, v_1)\| & = & \langle 0, \mathbf{map}^{\langle\langle\phi\rangle\rangle}(x.\|v_0\|_p, \|v_1\|_p) \rangle \\
\|\mathbf{let}(e_0, x.e_1)\| & = & \|e_0\|_c +_c \|e_1\|_c \langle \|e_0\|_p / x \rangle
\end{array}$$

Figure 3: Translation from source types and expressions to complexity types and expressions. Recall that $\|e\|_c = \pi_0 \|e\|$ and $\|e\|_p = \pi_1 \|e\|$.

We define the size-based complexity semantics as follows. The base cases for an inductive definition of (S^T, \leq_T) for every complexity type T consist of well-founded partial orders (S^Δ, \leq_Δ) for every datatype Δ in the signature, such that \leq_Δ is closed under arbitrary maximums (see below for a discussion). We define $\mathbf{N}^\infty = \mathbf{N} \cup \{\infty\}$, where \mathbf{N} is the natural numbers with the usual order and addition. We extend the order and addition to ∞ by $n \leq_{\mathbf{N}^\infty} \infty$ and $n + \infty = \infty + n = \infty + \infty = \infty$ for all $n \in \mathbf{N}$. For products and functions we define $S^{\mathbf{unit}} = \{*\}$ and $S^{T_0 \times T_1} = S^{T_0} \times S^{T_1}$ and $S^{T_0 \rightarrow T_1} = (S^{T_1})^{S^{T_0}}$, with the trivial, componentwise, and pointwise partial orders, respectively. Complexity types are interpreted into this type structure by setting $\llbracket \mathbf{C} \rrbracket = \mathbf{N}^\infty$ and $\llbracket T \rrbracket = S^T$ for each complexity type T .

Stating the conditions on programmer-defined size functions requires some auxiliary notions. For datatype $\Delta = \overline{C}$ of Φ_C , set $D^\Delta = \llbracket \Phi_{C_0}[\Delta] \rrbracket + \dots + \llbracket \Phi_{C_{n-1}}[\Delta] \rrbracket$, writing $\mathit{inj}_i : \llbracket \Phi_{C_i}[\Delta] \rrbracket \rightarrow D^\Delta$ for the i^{th} injection. Next, we define a function sz^Φ with domain $\llbracket \Phi[\Delta] \rrbracket$ (the semantic analogue of the argument type of a datatype constructor). $\mathit{sz}^\Phi(a)$ is intended to be the maximum of the values of type $\llbracket \Delta \rrbracket$ from which a is built using pairing and function application. We want to define sz^Φ by induction on Φ , computing the maximum at each step. To ignore values not of type $\llbracket \Delta \rrbracket$ we assume an element $\perp \notin S^\Delta$ that serves as an identity for \vee ; that is, we order $S^\Delta \cup \{\perp\}$ so that $\perp < a$ for all $a \in S^\Delta$. We define $\mathit{sz}^\Phi : \llbracket \Phi[\Delta] \rrbracket \rightarrow S^\Delta \cup \{\perp\}$ by induction on Φ as follows:

$$\begin{array}{l}
\mathit{sz}^\Phi(a) = a \\
\mathit{sz}^\Phi(a) = \perp \\
\mathit{sz}^{\Phi_0 \times \Phi_1}(a) = \mathit{sz}^{\Phi_0}(a) \vee \mathit{sz}^{\Phi_1}(a) \\
\mathit{sz}^{T \rightarrow \Phi}(f) = \bigvee_{a \in \llbracket T \rrbracket} \mathit{sz}^\Phi(f(a))
\end{array}$$

The key input to the size-based semantics is programmer-supplied size functions $\mathit{size}_\Delta : D^\Delta \rightarrow S^\Delta$ such that

$$\mathit{sz}^{\Phi C_i}(a) <_{S^\Delta \cup \{\perp\}} (\mathit{size}_\Delta \circ \mathit{inj}_i)(a)$$

size_Δ represents the programmer's notion of size for inductively-defined values. The only condition, which is used to interpret the recursor, is that the size of a value is strictly greater than the size of any of its substructures of the same type. For example, this condition permits interpreting the size of a list as its length or its total number of constructors, and the size of a tree as its number of nodes or its

height. Non-examples include defining the size of a list of natural numbers to be the number of successor constructors, and defining the size of all natural numbers to be a constant (though see Section 6.5 for a discussion of this latter possibility).

The interpretation of most terms is standard except for that of constructors and rec , which are given in Figure 4. We write $\mathit{map}^{\Phi, T_0, T_1}$ for semantic functions that mirror the definition of map , and we overload the notation C_i to stand for $\mathit{inj}_i : \llbracket \Phi_{C_i}[\delta] \rrbracket \rightarrow D^\delta$. The implementation of the recursors requires a bit of explanation, and is motivated by the goal to have $\|e\|$ bound the cost and potential of e . We expect that $\llbracket \mathbf{rec}^\delta(e, \overline{C \mapsto x.e_C}) \rrbracket$, which depends on $\llbracket \mathit{rec}^\delta(\|e\|_p, \overline{C \mapsto x.\|e_C\|}) \rrbracket$, should branch on $\llbracket \|e\|_p \rrbracket$, evaluating to the appropriate $\llbracket \|e_C\| \rrbracket$. However, $\llbracket \|e\|_p \rrbracket$ will be a semantic value of type S^δ , whereas to branch, we need a semantic value of type D^δ . Furthermore, $\llbracket \|e\|_p \rrbracket$ is an *upper bound* on the size of e , so $\llbracket \|e\|_p \rrbracket$ does not tell us the precise form of e , and so we cannot use $\llbracket \|e\|_p \rrbracket$ to predict which branch the evaluation of the source rec expression will follow. We solve these problems by introducing a semantic *case* function, and define the denotation of rec expressions by taking a maximum over the branches for all semantic values that are bounded by the upper bound $\llbracket \|e\|_p \rrbracket$. This is the source of the requirement that base-type potentials be closed under arbitrary maximums. Although this requirement seems rather strong, in most examples it seems easy to satisfy. In particular, we think of most datatype potentials (sizes) as being natural numbers, and so we satisfy the condition by interpreting them by \mathbf{N}^∞ .

The restriction on size_Δ ensures that the recursion used to interpret rec expressions descends along a well-founded partial order, and hence is well-defined. The maximum may end up being a maximum over all possible values, but this simply indicates that our interpretation fails to give us precise information.

We illustrate this semantics on some examples. In order to ease the notation, we will occasionally write syntactic expressions for the corresponding semantic values (in effect, dropping $\llbracket \cdot \rrbracket$). We also write the *case* function as a branch on constructors; for example, we write $\mathit{case}(t, \mathbf{Emp} \mapsto x.(1, 1) \mid \mathbf{Node} \mapsto \langle y, t_0, t_1 \rangle.e)$ for $\mathit{case}(t, \lambda x.(1, 1), \lambda \langle y, t_0, t_1 \rangle.e)$.

$$\begin{aligned}
\text{case}^\delta : D^\delta \times \prod_C (S^{\llbracket \Phi_C \rrbracket \delta} \rightarrow S^\tau) &\rightarrow S^\tau & \text{case}(Cx, (\dots, f_C, \dots)) &= f_C(x) \\
\llbracket Ce \rrbracket \xi &= \text{size}(C(\llbracket e \rrbracket \xi)) \\
\llbracket \text{rec}^\delta(E^\delta, \overline{C \mapsto x^{\phi_C[\delta \times \tau]}.E_C^\tau}) \rrbracket \xi &= \bigvee_{\text{size } z \leq \llbracket E \rrbracket \xi} \text{case}(z, (\dots, f_C, \dots)) \\
f_C(x) &= \llbracket E_C \rrbracket \xi \{x \mapsto \llbracket \text{map}^{\Phi_C}(w, \langle w, \text{rec}(w, \overline{C \mapsto x.E_C}) \rrbracket \xi, x) \rrbracket \xi, x\} \\
&= \llbracket E_C \rrbracket \xi \{x \mapsto \text{map}^{\Phi_C}(\lambda a. (a, \llbracket \text{rec}(w, \overline{C \mapsto x.E_C}) \rrbracket \xi \{w \mapsto a\}), x)\}
\end{aligned}$$

Figure 4: The interpretation of `rec` in the size-based semantics for the complexity language.

4.1 Booleans and Conditionals

In the source language we define booleans and their case construct:

```

datatype bool = True of unit | False of unit
case(ebool, e0τ, e1τ) = rec(e, True ↦ e0 | False ↦ e1)

```

(recall our convention on writing e_C for $x.e_C$ when $\phi_C = \text{unit}$). In the semantics of the complexity language, we interpret `bool` as a one-element set $\{1\}$, so `True` and `False` are indistinguishable by “size.” Our interpretation yields

$$\begin{aligned}
\llbracket \llbracket \text{case}(e, e_0, e_1) \rrbracket \rrbracket & \\
= \llbracket e \rrbracket_{c+c} & \\
\bigvee_{\text{size } b \leq \llbracket e \rrbracket_p} \text{case}(b, \text{True} \mapsto 1 + c \llbracket e_0 \rrbracket \mid \text{False} \mapsto 1 + c \llbracket e_1 \rrbracket) & \\
= \llbracket e \rrbracket_{c+c} (\text{case}(\text{True}, \text{True} \mapsto 1 + c \llbracket e_0 \rrbracket \mid \text{False} \mapsto 1 + c \llbracket e_1 \rrbracket) & \\
\vee \text{case}(\text{False}, \text{True} \mapsto 1 + c \llbracket e_0 \rrbracket \mid \text{False} \mapsto 1 + c \llbracket e_1 \rrbracket)) & \\
= (1 + \llbracket e \rrbracket_c) + c (\llbracket e_0 \rrbracket \vee \llbracket e_1 \rrbracket). &
\end{aligned}$$

In other words, if we cannot distinguish between `True` and `False` by size, then the interpretation of a conditional is just the maximum of its branches (with the additional cost of evaluating the test). This is precisely the interpretation used by Danner et al. (2013).

4.2 Tree Membership

Next we consider an example that shows that the “big” maximum used to interpret the recursor can typically be simplified to the recurrence that one expects to see. We analyze the cost of checking membership in an `int`-labeled tree. We write e_0 `orelse` e_1 as an abbreviation for `case` ($e_0, \text{True} \mapsto \text{True} \mid \text{False} \mapsto e_1$).

```

datatype tree = Emp of unit | Node of int × tree × tree
mem(t, x) = rec(t,
  Emp ↦ False
  Node ↦ ⟨y, ⟨t0, r0⟩, ⟨t1, r1⟩⟩.
    y = x orelse (force r0 orelse force r1)
)

```

For this example, we treat `int` (in the source and complexity languages) as a datatype with 2^{32} constructors where the equality test $x = y$ is implemented by a rather large case analysis. Let us define the size of a tree to be the number of nodes:

$$\begin{aligned}
\llbracket \text{tree} \rrbracket &= \mathbf{N}^\infty \\
D^{\text{tree}} &= \{*\} + \{1\} \times \mathbf{N}^\infty \times \mathbf{N}^\infty \\
\text{size}_{\text{tree}}(\text{Emp}) &= 0 \\
\text{size}_{\text{tree}}(\text{Node}(1, n_0, n_1)) &= 1 + n_0 + n_1
\end{aligned}$$

We would like to get the following recurrence for the cost of the `rec` expression when t has size n :

$$T(0) = 1 \quad T(n) = \bigvee_{n_0+n_1+1=n} 6 + T(n_0) + T(n_1)$$

($x = y$ requires an application and two `case` evaluations; each `orelse` evaluation costs 1; and we charge for the `rec` reduction).

Working through the interpretation yields $\llbracket \llbracket \text{mem}(t, x) \rrbracket \rrbracket_c = \llbracket t \rrbracket_c + g(\llbracket t \rrbracket_p) + 1$ where

$$\begin{aligned}
g(n) &= \llbracket \text{rec}(z, \text{Emp} \mapsto 1 \\
&\quad \text{Node} \mapsto \langle y, \langle t_0, r_0 \rangle, \langle t_1, r_1 \rangle \rangle. 6 + (r_0)_c + (r_1)_c \\
&\quad \rrbracket \{z \mapsto n\}.
\end{aligned}$$

We can calculate that $g(0) = 1$, and for $n > 0$:

$$\begin{aligned}
g(n) &= \bigvee_{\text{size } t \leq n} \text{case}(t, \\
&\quad \text{Emp} \mapsto 1 \\
&\quad \text{Node} \mapsto \langle y, n_0, n_1 \rangle. 6 + g(n_0) + g(n_1) \\
&= g(n-1) \vee \bigvee_{\text{size } t=n} \text{case}(t, \dots) \\
&= g(n-1) \vee \bigvee_{1+n_0+n_1=n} \text{case}(\text{Node}(1, n_0, n_1), \dots) \\
&= g(n-1) \vee \bigvee_{1+n_0+n_1=n} (6 + g(n_0) + g(n_1))
\end{aligned}$$

We now notice that when we take $n_0 = 0$ and $n_1 = n - 1$ we have

$$6 + g(n_0) + g(n_1) = 6 + g(0) + g(n-1) \geq g(n-1)$$

and hence

$$\begin{aligned}
g(n) &= g(n-1) \vee \bigvee_{1+n_0+n_1} (6 + g(n_0) + g(n_1)) \\
&= \bigvee_{1+n_0+n_1} (6 + g(n_0) + g(n_1))
\end{aligned}$$

which is precisely the recurrence we would expect.

4.3 Tree Map

Next, we consider an example that illustrates reasoning about higher-order functions and the benefits of choosing an appropriate notion of size. We analyze the cost of the `map` function for `nat`-labeled binary trees:

```

treemap(f, t) = rec(t,
  Emp ↦ Emp
  Node ↦ ⟨y, ⟨t0, r0⟩, ⟨t1, r1⟩⟩.
    Node(f(y), force r0, force r1)
)

```

Suppose the cost of evaluating f is monotone with respect to the size of its argument, where we define the size of a natural number n to be $1 + n$ (to count the zero constructor). The cost of evaluating `treemap`(f, t) should be bounded by $1 + n \cdot (1 + f(s)_c)$, where n is the number of nodes in t , s is the maximum size of all labels in t ,

and we write $f(s)_c$ for the cost of evaluating f on a natural number of size s (the map runs f on an input of size at most s for each of the n nodes, and takes an additional n steps to traverse the tree).

We take $\llbracket \text{tree} \rrbracket = \mathbf{N}^\infty \times \mathbf{N}^\infty$, where we think of the pair (n, s) as (number of nodes, maximum size of label), and use the mutual ordering on pairs $((n, s) < (n', s')$ iff $n \leq n'$ and $s < s'$ or $n < n'$ and $s \leq s')$. The size function is defined as follows:

$$\text{size}(\text{Emp}) = (0, 0)$$

$$\text{size}(\text{Node}(n, (n_0, s_0), (n_1, s_1))) = (1 + n_0 + n_1, \max\{n, s_0, s_1\}).$$

Let us write $g(m, s) = \llbracket \text{rec}(\dots) \rrbracket \{t \mapsto (m, s)\}$, so that $(\llbracket \text{treemap} \rrbracket (f, (m, s)))_c = g(m, s) + 1$. We now show that $g(m, s) \leq m(1 + f(s)_c)$ by induction:

$$\begin{aligned} g(m, s) &= \bigvee_{\text{size } z \leq (m, s)} \text{case}(z, \\ &\quad \text{Emp} \mapsto 1 \\ &\quad \text{Node} \mapsto \langle n, (n_0, s_0), (n_1, s_1) \rangle. \\ &\quad (1 + (f(n))_c + (g(n_0, s_0))_c + (g(n_1, s_1))_c) \\ &= 1 \vee \\ &\quad \bigvee_{\substack{1+n_0+n_1 \leq m \\ \max\{n, s_0, s_1\} \leq s}} (1 + f(n)_c + (g(n_0, s_0))_c + (g(n_1, s_1))_c) \\ &\leq \bigvee_{\substack{1+n_0+n_1 \leq m \\ \max\{n, s_0, s_1\} \leq s}} (1 + f(n)_c + \\ &\quad n_0 \cdot (1 + f(s_0)_c) + n_1 \cdot (1 + f(s_1)_c)) \\ &\leq \bigvee_{\substack{1+n_0+n_1 \leq m \\ \max\{n, s_0, s_1\} \leq s}} (1 + n_0 + n_1)(1 + f(\max\{n, s_0, s_1\})_c) \\ &\leq m \cdot (1 + f(s)_c). \end{aligned}$$

4.4 The Bounding Theorem for the Size-Based Semantics

The most basic correctness criterion for our technique is that a closed source program's operational cost is bounded by the cost component of the denotation of its complexity translation. However, to know that extracted *recurrences* are correct, it is not enough to consider closed programs; we also need to know that the potential of a function bounds that function's operational cost on all arguments, and so on at higher type. Thus, we use a logical relation. We first show a simplified case of the logical relation, where for this subsection only we do not allow datatype constructors to take functions as arguments (i.e., drop the $\tau \rightarrow \phi$ clause from constructor argument types ϕ). In Section 5, we consider the general case, which requires some non-trivial technical additions to the main definition.

DEFINITION 2.

1. Let e be a closed source language expression and a a semantic value. We write $e \sqsubseteq_\tau a$ to mean: if $e \downarrow^n v$, then
 - (a) $n \leq a_c$; and
 - (b) $v \sqsubseteq_\tau^{\text{val}} a_p$.
2. Let v be a source language value and a a semantic value. We define $v \sqsubseteq_\tau^{\text{val}} a$ by:
 - (a) $() \sqsubseteq_{\text{unit}}^{\text{val}} 1$.
 - (b) $\langle v_0, v_1 \rangle \sqsubseteq_{\tau_0 \times \tau_1}^{\text{val}} \langle a_0, a_1 \rangle$ if $v_i \sqsubseteq_{\tau_i}^{\text{val}} a_i$ for $i = 0, 1$.
 - (c) $\text{delay}(e) \sqsubseteq_{\text{susp } \tau}^{\text{val}} a$ if $e \sqsubseteq_\tau a$.

- (d) $C(v) \sqsubseteq_\delta^{\text{val}} a$ if there is a' such that $v \sqsubseteq_{\phi_C[\delta]}^{\text{val}} a'$ and $\text{size}(C(a')) \leq a$.¹
- (e) $\lambda x. e \sqsubseteq_{\sigma \rightarrow \tau}^{\text{val}} a$ if whenever $v \sqsubseteq_\sigma^{\text{val}} a'$, $e[v/x] \sqsubseteq_\tau a(a')$.

THEOREM 7 (Bounding theorem). *If $e : \tau$ in the source language, then $e \sqsubseteq_\tau \llbracket \llbracket e \rrbracket \rrbracket$.*

Rather than proving this bounding theorem directly, in Section 5 we identify syntactic constraints on the complexity language which allow the proof to be carried through (Theorem 12). Because the size-based semantics satisfies these syntactic constraints (see Section 6.1), we can prove that the logical relation defined in Section 5 implies the one defined above, giving Theorem 7 as a corollary.

5. The Syntactic Bounding Theorem

Rather than proving the bounding theorem for a particular model, such as the one from the previous section, we use a syntactic judgement $\Gamma \vdash E_0 \leq_\tau E_1$ to axiomatize the properties that are necessary to prove the theorem. The rules are in Figure 5; we omit typing premises from the figure, but formally each rule has sufficient premises to make the two terms have the indicated type. The first two rules state reflexivity and transitivity. The next rule (congruence) says that term contexts of a certain form (in the sequel, *congruence contexts*) are monotonic. The next three rules state the monoid laws for \mathbf{C} ; we write $E_0 = E_1$ to abbreviate two rules $E_0 \leq E_1$ and $E_1 \leq E_0$. The final three rules (which we call “step rules”) say that a β -redex is bigger than or equal to its reduct. The first five congruence contexts are the standard head elimination contexts used in logical relations arguments (principal arguments of elimination forms) and the next two say that $+$ is monotone.

These preorder rules are sufficient to prove the bounding theorem, and permit a variety of interpretations and extensions. If we impose no further rules, then $E_0 \leq E_1$ is basically weak head reduction from E_1 to E_0 (plus the monoid laws for \mathbf{C}). We can also add rules that identify elements of datatypes, in order to make those elements behave like sizes. For example, for lists of ints, we can say

$$\frac{}{E \leq \text{Cons}(_, E)} \quad \frac{}{\text{Cons}(E_1, E) \leq \text{Cons}(E_2, E)}$$

and extend the congruence contexts with $\text{Cons}(x, \mathcal{C})$. Then the second rule equates any two lists with the same number of elements, quotienting them to natural numbers, and the first rule orders these natural numbers by the usual less-than. Thus, considered up to \leq , lists are lengths.

Combining these rules with the ones used to prove the bounding theorem, the recursor for lists behaves like a monotonicization of the original recursion (like the \bigvee in the size-based complexity semantics). For example, for any specific list value $\text{Cons}(x, xs)$, by the usual step rule, we have

$$\frac{E_1[(x, xs, \text{rec}(xs, \text{Nil} \mapsto E_0, \text{Cons} \mapsto p.E_1))/p] \leq \text{rec}(\text{Cons}(x, xs), \text{Nil} \mapsto E_0, \text{Cons} \mapsto p.E_1)}{}{E_1[(x, xs, \text{rec}(xs, \text{Nil} \mapsto E_0, \text{Cons} \mapsto p.E_1))/p] \leq \text{rec}(\text{Cons}(x, xs), \text{Nil} \mapsto E_0, \text{Cons} \mapsto p.E_1)}$$

But we can derive $\text{Nil} \leq \text{Cons}(x, xs)$, so we also have

$$\begin{aligned} \text{rec}(\text{Nil}, \dots) &\leq \text{rec}(\text{Cons}(x, xs), \dots) && \text{by congruence} \\ E_0 &\leq \text{rec}(\text{Nil}, \text{Nil} \mapsto E_0, \text{Cons} \mapsto p.E_1) && \text{by the step rule} \\ E_0 &\leq \text{rec}(\text{Cons}(x, xs), \text{Nil} \mapsto E_0, \text{Cons} \mapsto p.E_1) && \text{by transitivity} \end{aligned}$$

and similarly for non-empty lists that are $\leq \text{Cons}(x, xs)$. Thus, when we quotient lists to their lengths, the congruence and step

¹ Our restriction on the form of ϕ_C allows us to conclude that this definition is well-founded, even though the type gets bigger in clause (2d), because we can treat the definition of $\sqsubseteq_\delta^{\text{val}}$ as an inner induction on the values. Allowing datatype constructors to take function arguments complicates the situation, and in Section 5 we must define a more general relation.

rules for `rec` (used to prove the bounding theorem) imply that the recursor is bigger than all of the branches for all smaller lists.

In Section 4, we used reasoning in the size-based semantics to massage the recurrence extracted from a program into a recognizable and solvable form. In future work, we plan to investigate how to do this massaging within the syntax of complexity language, using the rules we have just discussed and others. For example, while a recursion bounds what it steps to on all smaller values, we do not yet have a rule stating that it is a least upper bound. Here, we lay a foundation for this by proving the bounding theorem for the small set of rules in Figure 5.

5.1 The Bounding Relation

First, we extend Definition 2 to arbitrary datatypes. Fix a signature ψ . We will mutually define the following relations:

1. $e \sqsubseteq_{\tau} E$, where $\emptyset \vdash_{\psi} e : \tau$ and $\emptyset \vdash_{\|\psi\|} E : \|\tau\|$.
2. $v \sqsubseteq_{\tau}^{\text{val}} E$, where $\emptyset \vdash_{\psi} v : \tau$ and $\emptyset \vdash_{\|\psi\|} E : \langle\langle\tau\rangle\rangle$.
3. $v \sqsubseteq_{\phi, R}^{\text{val}} E$, where $\emptyset \vdash_{\psi} v : \phi[\delta]$ and $\emptyset \vdash_{\|\psi\|} E : \langle\langle\phi\rangle\rangle[\delta]$.
4. $e \sqsubseteq_{\phi, R} E$, where $\emptyset \vdash_{\psi} e : \phi[\delta]$ and $\emptyset \vdash_{\|\psi\|} E : \|\phi\|[\delta]$

In (3) and (4), $R(\emptyset \vdash_{\psi} v : \delta, \emptyset \vdash_{\|\psi\|} E : \delta)$, is any relation; these parts interpret strictly positive functors as relation transformers.

The definition is by induction on τ and ϕ . For datatypes, the signature well-formedness relation $\psi \text{ sig}$ ensures that datatypes are ordered, where later ones can refer to earlier ones, but not vice versa. Therefore, we could “inline” all datatype declarations: rather than naming datatypes, we could replace each datatype name δ by an inductive type $\mu[C \text{ of } \phi]$. The logical relation is defined using the subterm ordering for this “inlined” syntax. In addition to the usual subterm ordering for types τ and functors ϕ , we have that datatypes that occur earlier in ψ are smaller than later ones, and if $C : (\phi \rightarrow \delta) \in \psi$, then ϕ is smaller than δ .

DEFINITION 3.

1. We write $e \sqsubseteq_{\tau} E$ to mean: if $e \downarrow^n v$, then
 - $n \leq E_c$; and
 - $v \sqsubseteq_{\tau}^{\text{val}} E_p$.
2. We write $v \sqsubseteq_{\tau}^{\text{val}} E$ to mean:
 - $v \sqsubseteq_{\text{unit}}^{\text{val}} E$ is always true.
 - $\langle v_1, v_2 \rangle \sqsubseteq_{\tau_1 \times \tau_2}^{\text{val}} E$ iff $v_1 \sqsubseteq_{\tau_1}^{\text{val}} \pi_0 E$ and $v_2 \sqsubseteq_{\tau_2}^{\text{val}} \pi_1 E$.
 - $\text{delay}(e) \sqsubseteq_{\text{susp } \tau}^{\text{val}} E$ iff $e \sqsubseteq_{\tau} E$.
 - $v \sqsubseteq_{\delta}^{\text{val}} E$ is inductively defined by
$$\frac{C : (\phi \rightarrow \delta) \in \psi \quad v \sqsubseteq_{\phi, -\sqsubseteq_{\delta}^{\text{val}}}^{\text{val}} E' \quad C E' \leq \delta E}{C v \sqsubseteq_{\delta}^{\text{val}} E}$$
- $\lambda x. e \sqsubseteq_{\tau_1 \rightarrow \tau_2}^{\text{val}} E$ iff (for all v_1 and E_1 , if $v_1 \sqsubseteq_{\tau_1}^{\text{val}} E_1$ then $e[v_1/x] \sqsubseteq_{\tau_2} E E_1$).
3. We write $v \sqsubseteq_{\phi, R}^{\text{val}} E_p$ to mean:
 - $v \sqsubseteq_{t, R}^{\text{val}} E$ if $R(v, E)$.
 - $v \sqsubseteq_{\tau, R}^{\text{val}} E$ if $v \sqsubseteq_{\tau}^{\text{val}} E$ (t not free in τ).
 - $\langle v, v' \rangle \sqsubseteq_{\phi \times \phi', R}^{\text{val}} E$ if $v \sqsubseteq_{\phi, R}^{\text{val}} \pi_0 E$ and $v' \sqsubseteq_{\phi', R}^{\text{val}} \pi_1 E$.
 - $\lambda x. e_1 \sqsubseteq_{\tau \rightarrow \phi, R}^{\text{val}} E_1$ if for all v and E , if $v \sqsubseteq_{\tau}^{\text{val}} E$, then $e_1[v/x] \sqsubseteq_{\phi, R} (E_1 E)$.
4. We write $e \sqsubseteq_{\phi, R} E$ to mean: if $e \downarrow^n v$, then
 - $n \leq E_c$; and
 - $v \sqsubseteq_{\phi, R}^{\text{val}} E_p$.

The inner inductive definition of $v \sqsubseteq_{\delta}^{\text{val}} E$ makes sense because R occurs strictly positively in $-\sqsubseteq_{\phi, R}^{\text{val}}-$, and because (by signature formation) δ cannot occur in ϕ , so $-\sqsubseteq_{\delta}^{\text{val}}-$ does not occur

elsewhere in $-\sqsubseteq_{\phi, R}^{\text{val}}-$. The relation on open terms considers all closed instances:

5. For a source substitution $\theta : \gamma$ and complexity substitution $\Theta : \Gamma$, we write $\theta \sqsubseteq_{\gamma}^{\text{sub}} \Theta$ to mean that for all $(x : \tau) \in \gamma$, $\theta(x) \sqsubseteq_{\tau}^{\text{val}} \Theta(x)$.
6. For $\gamma \vdash e : \tau$ and $\Gamma \vdash E : \|\tau\|$, we write $e \sqsubseteq_{\tau} E$ to mean that for all $\theta : \gamma$ and $\Theta : \Gamma$, if $\theta \sqsubseteq_{\gamma}^{\text{sub}} \Theta$, then $e[\theta] \sqsubseteq_{\tau} E[\Theta]$.

We write $\mathcal{E} :: \mathcal{J}$ to mean that \mathcal{E} is a derivation of any of the judgements just described. Because the relation for function types is a function between relations, derivations are infinitely-branching trees. A *subderivation* of such an \mathcal{E} is any subtree of \mathcal{E} , which includes any application of an \rightarrow -type judgement. For example, if $\mathcal{E}_1 :: \lambda x. e_1 \sqsubseteq_{\tau \rightarrow \phi, R}^{\text{val}} E_1$ and $\mathcal{E} :: v \sqsubseteq_{\tau}^{\text{val}} E$, then the derivation of $e_1[v/x] \sqsubseteq_{\phi, R}^{\text{val}} E_1 E$ is a subderivation of \mathcal{E}_1 .

Next, we establish some basic properties of the relation:

LEMMA 8 (Weakening).

1. If $e \sqsubseteq_{\tau} E$ and $E \leq_{\|\tau\|} E'$ then $e \sqsubseteq_{\tau} E'$.
2. If $v \sqsubseteq_{\tau}^{\text{val}} E$ and $E \leq_{\langle\langle\tau\rangle\rangle} E'$ then $v \sqsubseteq_{\tau}^{\text{val}} E'$.

Proof. Both clauses are proved simultaneously by induction on τ , using congruence for $\pi_0 []$, $\pi_1 []$ and $[] E$. See the full paper for details. \square

LEMMA 9 (Compositionality).

1. $e \sqsubseteq_{\phi, -\sqsubseteq_{\tau}^{\text{val}}} E$ iff $e \sqsubseteq_{\phi[\tau]} E$.
2. $v \sqsubseteq_{\phi, -\sqsubseteq_{\tau}^{\text{val}}}^{\text{val}} E$ iff $v \sqsubseteq_{\phi[\tau]}^{\text{val}} E$.

Proof. (1) follows by post-composing with (2), and (2) follows by induction on ϕ . See the full paper for details. \square

5.2 The Fundamental Theorem

First we state two lemmas which say that, when applied to related arguments, source-language `map` is bounded by complexity-language `map`, and that source-language `rec` is bounded by complexity-language `rec`.

LEMMA 10 (Map). *Suppose:*

1. $x : \tau_0 \vdash v_1 : \tau_1$ and $\emptyset \vdash v_0 : \phi[\tau_0]$;
2. $x : \langle\langle\tau_0\rangle\rangle \vdash E_1 : \langle\langle\tau_1\rangle\rangle$ and $\emptyset \vdash E_0 : \langle\langle\phi\rangle\rangle[\langle\langle\tau_0\rangle\rangle]$;
3. $\mathcal{E} :: v_0 \sqsubseteq_{\phi, -\sqsubseteq_{\tau_0}^{\text{val}}}^{\text{val}} E_0$;
4. Whenever \mathcal{E}' is a subderivation of \mathcal{E} such that $\mathcal{E}' :: v'_0 \sqsubseteq_{\tau_0}^{\text{val}} E'_0$, $v_1[v'_0/x] \sqsubseteq_{\tau_0}^{\text{val}} E_1[E'_0/x]$; and
5. $\text{map}^{\phi}(x.v_1, v_0) \downarrow^n v$.

Then $n = 0$ and $v \sqsubseteq_{\phi[\tau_0]}^{\text{val}} \text{map}^{\langle\langle\phi\rangle\rangle}(x.E_1, E_0)$.²

Proof. The proof is by induction on ϕ . Lemma 3 shows that $n = 0$. Omitted cases are in the full paper.

CASE: $\phi = \tau \rightarrow \phi_0$. Then $v_0 = \lambda y. e_0$ and \mathcal{E} proves that for all $v' \sqsubseteq_{\tau}^{\text{val}} E'$, $e_0[v'/y] \sqsubseteq_{\phi_0, -\sqsubseteq_{\tau_0}^{\text{val}}} E_0(E')$. Since $v_0 = \lambda y. e_0$, $v = \lambda y. \text{let}(e_0, z. \text{map}^{\phi}(x.v_1, z))$, so we must show that $\lambda y. \text{let}(e_0, z. \text{map}^{\phi}(x.v_1, z)) \sqsubseteq_{\tau \rightarrow \phi_0[\tau_0]}^{\text{val}} \text{map}^{\langle\langle\tau \rightarrow \phi_0\rangle\rangle}(x.E_1, E_0)$. To do so, suppose $w \sqsubseteq_{\tau}^{\text{val}} F$; we must show that

$$\text{let}(e_0[w/y], z. \text{map}^{\phi}(x.v_1, z)) \sqsubseteq_{\phi_0[\tau_0]}^{\text{val}} \text{map}^{\|\phi_0\|}(x.E_1, E_0(F)). \quad (*)$$

²We could have said $\text{map}^{\phi}(x.v_1, v_0) \sqsubseteq_{\phi[\tau_0]} \langle 0, \text{map}^{\langle\langle\phi\rangle\rangle}(x.E_1, E_0) \rangle$ but this version of the lemma avoids needing the symmetric copy of the step rule for pairs.

$$\begin{array}{c}
C ::= [] \mid \pi_0 C \mid \pi_1 C \mid C E \mid \text{rec}(C, \overline{C \mapsto x.E_C}) \mid C + E \mid E + C \\
\frac{}{\Gamma \vdash E \leq_T E} \quad \frac{\Gamma \vdash E_0 \leq_T E_1 \quad \Gamma \vdash E_1 \leq_T E_2}{\Gamma \vdash E_0 \leq_T E_2} \quad \frac{\Gamma, x : T' \vdash C[x] : T \quad \Gamma \vdash E_0 \leq_{T'} E_1}{\Gamma \vdash C[E_0] \leq_T C[E_1]} \text{ (congruence)} \\
\frac{}{\Gamma \vdash 0 + E =_C E} \quad \frac{}{\Gamma \vdash E + 0 =_C E} \quad \frac{}{\Gamma \vdash (E_0 + E_1) + E_2 =_C E_0 + (E_1 + E_2)} \\
\frac{\Gamma \vdash E_0[E_1/x] \leq_T (\lambda x.E_0)E_1 \quad \Gamma \vdash E_i \leq_{T_i} \pi_i \langle E_0, E_1 \rangle}{C : (\Phi \rightarrow \Delta) \in \Psi} \\
\frac{}{\Gamma \vdash E_C[\text{map}^\Phi(y.\langle y, \text{rec}(y, \overline{C \mapsto x.E_C}) \rangle), E_0]/x] \leq_T \text{rec}^\Delta(CE_0, \overline{C \mapsto x.E_C})}
\end{array}$$

Figure 5: Congruence contexts and the preorder judgement

Suppose

$$\frac{e_0[w/y] \downarrow^{n_0} w_0 \quad \text{map}^{\phi_0}(x.v_1, w_0) \downarrow^{n_1} v'}{\text{let}(e_0[w/y], z.\text{map}^{\phi_0}(x.v_1, z)) \downarrow^{n_0+n_1} v'}$$

Since $w \sqsubseteq_{\tau}^{\text{val}} F$, we have that \mathcal{E} derives $e_0[w/y] \sqsubseteq_{\phi_0, -\sqsubseteq_{\tau_0}^{\text{val}}} E_0(F)$ and hence we have a subderivation \mathcal{E}_0 of \mathcal{E} such that $\mathcal{E}_0 :: w_0 \sqsubseteq_{\phi_0, -\sqsubseteq_{\tau_0}^{\text{val}}} (E_0(F))_p$. We now verify that (4) holds for \mathcal{E}_0 so that we can apply the induction hypothesis to $\text{map}^{\phi}(x.v_1, w_0)$. So suppose that \mathcal{E}'_0 is a subderivation of \mathcal{E}_0 such that $\mathcal{E}'_0 :: w'_0 \sqsubseteq_{\tau_0}^{\text{val}} F'_0$. We need to show that $v_1[w'_0/x] \sqsubseteq_{\tau_0}^{\text{val}} E_1[F'_0/x]$, and to do so it suffices to note that \mathcal{E}'_0 is a subderivation of \mathcal{E}_0 , which in turn is a subderivation of \mathcal{E} .

We can now apply the induction hypothesis to conclude that $n_1 = 0$ and so:

$$\begin{aligned}
n_0 + n_1 &= n_0 \leq (E_0 F)_c = (\text{map}^{\|\phi\|}(x.E_1, E_0 F))_c \\
v' \sqsubseteq_{\phi[\tau_0]}^{\text{val}} \text{map}^{\langle\phi\rangle}(x.E_1, (E_0 F)_p) &= (\text{map}^{\|\phi\|}(x.E_1, E_0 F))_p.
\end{aligned}$$

Using β for pairs, these are the two conditions that must be verified to show (*), so this completes the proof. \square

LEMMA 11 (Recursor). Fix a datatype declaration $\text{datatype } \delta = \overline{C} \text{ of } \phi$. If $v \sqsubseteq_{\delta}^{\text{val}} E$ and for all C , $e_C \sqsubseteq_{\phi_C[\delta \times \text{susp } \tau]} E_C$, then $\text{rec}(v, \overline{C \mapsto x.e_C}) \sqsubseteq \text{rec}(E, \overline{C \mapsto x.1 +_c E_C})$

Proof. By induction on $v \sqsubseteq_{\delta}^{\text{val}} E$. The only case is

$$\frac{C : (\phi \rightarrow \delta) \in \psi \quad v' \sqsubseteq_{\phi, -\sqsubseteq_{\delta}^{\text{val}}} E' \quad C E' \leq_{\delta} E}{C v' \sqsubseteq_{\delta}^{\text{val}} E} \quad (\dagger)$$

Assume $\text{rec}(C v', \overline{C \mapsto x.e_C})$ evaluates. Then by inversion and Lemma 2 it was by

$$\frac{\begin{array}{l} C v' \downarrow^0 C v' \\ \text{map}^{\phi}(y.\langle y, \text{delay}(\text{rec}(y, \overline{C \mapsto x.e_C})) \rangle, v') \downarrow^0 v'' \\ e_C[v''/x] \downarrow^{n_2} v \end{array}}{\text{rec}(C v', \overline{C \mapsto x.e_C}) \downarrow^{0+1+n_2} v} \quad (*)$$

Using the premise that $C E' \leq_{\delta} E$ from (\dagger), β for datatypes, and congruence, we note that

$$\begin{aligned}
&\text{rec}(E, \overline{C \mapsto x.1 +_c E_C}) \\
&\geq \text{rec}(C E', \overline{C \mapsto x.1 +_c E_C}) \\
&\geq 1 +_c E_C[\text{map}^{\langle\phi\rangle}(y.\langle y, \text{rec}(y, \overline{C \mapsto x.1 +_c E_C}) \rangle, E')/x]
\end{aligned}$$

Let us write E^* for $\text{map}^{\langle\phi\rangle}(y.\langle y, \text{rec}(y, \overline{C \mapsto x.1 +_c E_C}) \rangle, E')$. Thus by congruence, transitivity, weakening, and β for pairs, it suffices to show

$$\begin{aligned}
1 + n_2 &\leq 1 + E_C[E^*/x]_c \\
v &\sqsubseteq^{\text{val}} (E_C[E^*/x])_p
\end{aligned}$$

By congruence for $+$, for the first goal it suffices to show $n_2 \leq E_C[E^*/x]_c$. Thus, if we can show $e_C[v''/x] \sqsubseteq E_C[E^*/x]$, then applying it to the third evaluation premise of (*) gives the result. We can use our assumption that $e_C \sqsubseteq E_C$, as long as we show $v'' \sqsubseteq^{\text{val}} E^*$. To do so, we use Lemma 10 applied to the second evaluation premise of (*) with

$$\begin{aligned}
v_1 &= v' & v &= y.\langle y, \text{delay}(\text{rec}(y, \overline{C \mapsto x.e_C})) \rangle \\
E_1 &= E' & E &= y.\langle y, \text{rec}(y, \overline{C \mapsto x.1 +_c E_C}) \rangle
\end{aligned}$$

We have $\mathcal{E} :: v' \sqsubseteq_{\phi, -\sqsubseteq_{\delta}^{\text{val}}} E'$ from the second premise of (\dagger). Thus, to finish calling the theorem, we need to show that for all R -position subderivations of \mathcal{E} deriving $v'_1 \sqsubseteq_{\delta}^{\text{val}} E'_1$,

$$\begin{aligned}
\langle v'_1, \text{delay}(\text{rec}(v'_1, \overline{C \mapsto x.e_C})) \rangle &\sqsubseteq_{\delta \times \text{susp } \tau}^{\text{val}} \\
\langle E'_1, \text{rec}(E'_1, \overline{C \mapsto x.1 +_c E_C}) \rangle &
\end{aligned}$$

By definition of value bounding at product types, weakening and β for pairs, it suffices to show

$$v'_1 \sqsubseteq_{\delta}^{\text{val}} E'_1$$

$$\text{delay}(\text{rec}(v'_1, \overline{C \mapsto x.e_C})) \sqsubseteq_{\text{susp } \tau}^{\text{val}} \text{rec}(E'_1, \overline{C \mapsto x.1 +_c E_C})$$

The former we have, and for the latter by definition it suffices to show

$$\text{rec}(v'_1, \overline{C \mapsto x.e_C}) \sqsubseteq_{\tau} \text{rec}(E'_1, \overline{C \mapsto x.1 +_c E_C})$$

Because $v'_1 \sqsubseteq_{\delta}^{\text{val}} E'_1$ is an R -subderivation of $v' \sqsubseteq_{\phi, -\sqsubseteq_{\delta}^{\text{val}}} E'$, and therefore a strict subderivation of $C v' \sqsubseteq_{\delta}^{\text{val}} E$, we can use the inductive hypothesis on it, which gives exactly what we needed to show. \square

THEOREM 12 (Bounding Theorem). If $\gamma \vdash e : \tau$, then $e \sqsubseteq_{\tau} \|e\|$.

Proof. By induction on the derivation of $\gamma \vdash e : \tau$. In each case we state the last line of the derivation, taking as given the premises of the typing rules in Figure 1. Omitted cases are in the full paper.

CASE: $\gamma \vdash \text{rec}(e, \overline{C \mapsto x.e_C}) : \tau$. We need to show

$$\text{rec}(e[\theta], \overline{C \mapsto x.e_C[\theta, x/x]}) \sqsubseteq \langle E_c + (E_r)_c, (E_r)_p \rangle$$

where $E = \|e\|[\Theta]$ and $E_r = \text{rec}(E_p, \overline{C \mapsto x.(1 +_c \|e_C\|[\Theta, x/x])}$.
Suppose

$$\frac{e[\theta] \downarrow^{n_0} C v_0 \quad \text{map}^{\phi_C}(y.\langle y, \text{delay}(\text{rec}(y, \overline{C \mapsto x.e_C[\theta, x/x]})) \rangle, v_0) \downarrow^0 v_1 \quad e_C[\theta, x/v_1] \downarrow^{n_2} v}{\text{rec}(e[\theta], \overline{C \mapsto x.e_C[\theta, x/x]}) \downarrow^{1+n_0+n_2} v}$$

By the induction hypothesis $e[\theta] \sqsubseteq E$, so $n_0 \leq E_c$ and $C v_0 \sqsubseteq^{\text{val}} E_p$. By Lemma 2 we can derive

$$\frac{C v_0 \downarrow^0 C v_0 \quad \text{map}^{\phi_C}(y.\langle y, \text{delay}(\text{rec}(y, \overline{C \mapsto x.e_C[\theta, x/x]})) \rangle, v_0) \downarrow^0 v_1 \quad e_C[\theta, x/v_1] \downarrow^{n_2} v}{\text{rec}(C v_0, \overline{C \mapsto x.e_C[\theta, x/x]}) \downarrow^{1+n_2} v}$$

So by Lemma 11 we have that $1 + n_2 \leq (E_r)_c$ and $v \sqsubseteq^{\text{val}} (E_r)_p$. Putting these together, we have what we needed to show:

$$1 + n_0 + n_2 \leq E_c + (E_r)_c \quad v \sqsubseteq^{\text{val}} (E_r)_p$$

CASE: $\gamma \vdash \text{map}^{\phi}(x.v_1, v_0) : \phi[\tau_1]$. Because v_1 is a sub-syntactic-class of e , we can upcast it and apply $\|v_1\|$ to it, producing a complexity expression. We must show that

$$\text{map}^{\phi}(x.v_1[\theta, x/x], v_0[\theta]) \sqsubseteq (0, \text{map}^{\langle \phi \rangle}(x.\|v_1\|[\Theta, x/x]_p, \|v_0\|[\Theta]_p)),$$

so suppose $\text{map}^{\phi}(x.v_1[\theta, x/x], v_0[\theta]) \downarrow^n v$. By transitivity/weakening with β for pairs, it suffices to show:

$$n \leq 0 \quad v \sqsubseteq^{\text{val}} \text{map}^{\langle \phi \rangle}(\|v_1\|[\Theta, x/x]_p, \|v_0\|[\Theta]_p) \quad (*)$$

We will apply Lemma 10 to $\text{map}^{\phi}(x.v_1[\theta, x/x], v_0[\theta]) \downarrow^n v$ with

$$\begin{aligned} v_0 &= v_0[\theta] & v_1 &= v_1[\theta, x/x] \\ E_0 &= \|v_0\|[\Theta]_p & E_1 &= \|v_1\|[\Theta, x/x]_p \end{aligned}$$

To establish condition (3) we apply the IH to v_0 to conclude that $v_0[\theta] \sqsubseteq_{\phi[\tau_0]} \|v_0\|[\Theta]$. Since $v_0[\theta]$ is a value, by Lemma 2, it evaluates to itself. Therefore $v_0[\theta] \sqsubseteq_{\phi[\tau_0]}^{\text{val}} \|v_0\|[\Theta]_p$ and so by Lemma 9, $v_0[\theta] \sqsubseteq_{\phi, -\sqsubseteq^{\text{val}}_0}^{\text{val}} \|v_0\|[\Theta]_p$.

To establish condition (4), assume $v'_0 \sqsubseteq_{\tau'_0}^{\text{val}} E'_0$ (which is an R -subderivation of the above, but we won't use this fact). Using the substitution lemmas we need to show $v_1[\theta, v'_0/x] \sqsubseteq^{\text{val}} \|v_1\|[\Theta, E'_0/x]_p$. Since $\theta, v'_0/x \sqsubseteq^{\text{sub}} \Theta, E'_0/x$, the IH on v_1 gives $v_1[\theta, v'_0/x] \sqsubseteq \|v_1\|[\Theta, E'_0/x]$ and since $v_1[\theta, v'_0/x]$ is a value, it evaluates to itself, so $v_1[\theta, v'_0/x] \sqsubseteq^{\text{val}} \|v_1\|[\Theta, E'_0/x]_p$ as we needed to show.

Now we apply Lemma 10 to conclude (*). \square

6. Models of the Complexity Language

A model of the complexity language consists of an interpretation of types as preorders, and of terms as maps between elements of those preorders, validating the rules of Figure 5. The congruence contexts C , but not all terms, need to be monotone maps.

6.1 The Size-Based Complexity Semantics

We showed in Section 4 that the size-based semantics interprets the syntax of the complexity language; it is also a model of the preorder rules of Figure 5. Congruence is established by induction on C ; we

do not need programmer-defined size functions to be monotonic, because there is no congruence context for datatype constructors.

$$\begin{aligned} & \llbracket \text{rec}(C E_0, \overline{x \mapsto E_C}) \rrbracket \xi \\ &= \bigvee_{\text{size } z \leq \text{size}(C \llbracket E_0 \rrbracket \xi)} \text{case}(z, (\dots, f_C, \dots)) \\ &\geq \text{case}(C \llbracket E_0 \rrbracket \xi, (\dots, f_C, \dots)) \\ &= \llbracket E_C \rrbracket \xi \{x \mapsto \llbracket \text{map}^{\phi_C}(w.\langle w, \text{rec}(w, \overline{x \mapsto E_C}) \rangle, E_0) \rrbracket \xi\}. \end{aligned}$$

Verification of the pre-order axioms is straightforward. Therefore, Theorem 7 is a corollary of Theorem 12.

6.2 Infinite-Width Trees

Infinite-width trees can be defined by the declaration

$$\text{datatype tree} = E \text{ of unit} \mid N \text{ of int} \times (\text{nat} \rightarrow \text{tree})$$

Though every branch in such a tree is of finite length, the height of a tree is in general not a finite natural number. However, the size-based semantics adapts easily to interpret `tree` by a suitably large infinite successor ordinal, and then defining $\text{size}(N(x, f)) = \bigvee_{y \in \llbracket \text{nat} \rrbracket} f(y) + 1$.

6.3 A Semantics Without Arbitrary Maximums

The language studied in Danner et al. (2013) can be viewed as a specific signature in the present language. Their language has a type of booleans, a type `int` of fixed-size integers, and a type `list` of integer lists. As in Example 4.2, we can treat `int` and `bool` as enumerated datatypes with unit-cost operations. The `list` type is defined as a datatype and its case and fold operators are easily defined using `rec`.

For this specific signature, we can give a semantics of the complexity language in which we interpret `list` by \mathbf{N} instead of \mathbf{N}^∞ . Set $\llbracket \text{Nil} \rrbracket \xi = 0$ and $\llbracket \text{Cons}(E_0, E_1) \rrbracket \xi = \llbracket E_1 \rrbracket \xi + 1$. Define a semantic primitive recursion operator by $\text{rec}(0, a, f) = a$ and $\text{rec}(n+1, a, f) = a \vee f(n, \text{rec}(n, a, f))$. Finally, set

$$\begin{aligned} & \llbracket \text{rec}(E) \rrbracket \xi = \\ & \text{rec}(\llbracket E \rrbracket \xi, \llbracket E_{\text{Nil}} \rrbracket \xi, \lambda n, w. \llbracket E_{\text{Cons}} \rrbracket \xi \{x, xs, r \mapsto 1, n, w\}). \end{aligned}$$

where $\text{rec}(E) = \text{rec}(E, \text{Nil} \mapsto E_{\text{Nil}}, \text{Cons} \mapsto \langle x, \langle xs, r \rangle \rangle.E_{\text{Cons}})$. Verifying the preorder rules from Figure 5 is straightforward in all cases except the last; we verify the `Cons` case as follows:

$$\begin{aligned} & \llbracket \text{rec}(\text{Cons}(E_0, E_1)) \rrbracket \xi \\ &= (\llbracket E_{\text{Nil}} \rrbracket \xi \{x \mapsto 1\}) \vee \\ & \quad (\llbracket E_{\text{Cons}} \rrbracket \xi \{x, xs, r \mapsto 1, \llbracket E_1 \rrbracket \xi, \text{rec}(\llbracket E_1 \rrbracket \xi, \dots)\}) \\ &\geq \llbracket E_{\text{Cons}} \rrbracket \xi \{x, xs, r \mapsto 1, \llbracket E_1 \rrbracket \xi, \text{rec}(\llbracket E_1 \rrbracket \xi, \dots)\} \\ &= \llbracket E_{\text{Cons}}[E_0, (E_1, \text{rec}(E_1))/x, \langle xs, r \rangle] \rrbracket \xi \\ &= \llbracket E_{\text{Cons}}[E_0, \text{map}(y.\langle y, \text{rec}(y) \rangle, \langle E_0, E_1 \rangle)/x, xs, r] \rrbracket \xi. \end{aligned}$$

6.4 Exact Costs

If we wish to reason about exact costs, we can symmetrize the inequalities in Figure 5 into equalities, and add congruence for all contexts, which makes the $E_0 \leq E_1$ judgement into a standard notion of definitional equality. Then we can take the term model in the usual way, interpreting each type as a set of terms quotiented by this definitional equality. In this interpretation $\|e\|_c$ is a recurrence that gives the exact cost of evaluating e , but reasoning about such a recurrence involves reasoning about all of the details of the program.

6.5 Infinite Costs

Next, we consider a size-based model in which we drop the ‘‘increasing’’ requirement on the *size* functions from Section 4. Rather

than requiring a well-founded partial order for each datatype, we require an arbitrary partial order (S^τ, \leq_τ) which we also interpret as a flat CPO (we do not require the interpretation of non-datatypes to be CPOs). The interpretation of rec expressions is then in terms of a general fixpoint operator. Define $\infty = \bigvee S^\Delta$ and identify ∞ with the bottom element of the CPO ordering. In this setting it may be that the interpretation of a rec expression does not terminate and hence, by our identification, evaluates to ∞ . This turns out to be exactly the right behavior, as we can see in the following example.

Take the standard inductive definition of `nat` and interpret `nat` as some one-element set $\{1\}$ in the complexity language. Now compute the interpretation of the identity function:

$$\begin{aligned} & \llbracket \text{rec}(y, \text{Zero} \mapsto \text{Zero}, \text{Succ} \mapsto x.\text{Succ } x) \rrbracket \\ &= e(1) \\ &= \bigvee_{\text{size } z \leq 1} \text{case}(z, \text{Zero} \mapsto (0, 1) \mid \text{Succ} \mapsto \langle x, r \rangle.1 + e_c(x)) \end{aligned}$$

where

$$e(x) = \text{rec}(x, \text{Zero} \mapsto (0, 1) \mid \text{Succ} \mapsto \langle x, r \rangle.(1 + r_c, 1))$$

Since $\text{size}(\text{Succ}(1)) = 1 \leq 1$, one of the *case* expressions in the maximum is $e_c(1)$. In other words, we have a non-terminating recursion in computing the complexity. We conclude $\llbracket \text{rec}(\dots) \rrbracket_c = \infty$; in other words, we can draw no useful conclusion about the cost of this expression. What we have done in this example is to declare that we cannot distinguish values of type `nat` by size, and then we attempt to compute the cost of a recursive function on `nats` in terms of the size of the recursion argument. The bound given by the bounding theorem is correct, just not useful; it does not even tell us that the computation terminates.

7. Related Work

There is a reasonably extensive literature over the last several decades on (semi-)automatically constructing resource bounds from source code. The first work concerns itself with first-order programs. Wegbreit (1975) describes a system for analyzing simple Lisp programs that produces closed forms that bound running time. An interesting aspect of this system is that it is possible to describe probability distributions on the input domain and the generated bounds incorporate this information. Rosendahl (1989) proposes a system based on step-counting functions and abstract interpretation for a first-order subset of Lisp. More recently the COSTA project (see, e.g., Albert et al. (2012)) has focused on automatically computing cost relations for imperative languages (actually, bytecode) and solving them (more on that in the next section). Debray and Lin (1993) develop a system for analyzing logic programs and Navas et al. (2007) extend it to handle user-defined resources.

The Resource Aware ML project (RAML) takes a different approach to the one we have described here, one based on type assignment. Jost et al. (2010) describe a formalism that automatically infers linear resource bounds for higher-order programs, provided that the input program does in fact have a linear resource cost. Hoffmann and Hofmann (2010) and Hoffmann et al. (2012) extend this work to handle polynomial bounds, though for first-order programs only, and Hoffmann and Shao (2015) extend it to parallel programs. RAML uses a source language that is similar to ours, but in which the types are annotated with variables corresponding to resource usage. Type inference in the annotated system comes down to solving a set of constraints among these variables. A very nice feature of this work is that it handles cases in which amortized analysis is typically employed to establish tight bounds, while our approach can only conclude (worst-case) bounds.

Danielsson (2003) uses an annotated monad (similar to $\mathbf{C} \times -$) to track running time in a dependently typed language, where *size*

reasoning can be done via types. He emphasizes reasoning about amortized cost of lazy programs. He relies on explicit annotation of the program, which our complexity translation inserts automatically, and his correctness theorem is for closed programs, whereas we use a logical relation to validate extracted recurrences.

We now turn to work that is closest in spirit to ours, focusing on those aspects related to analysis of higher-order languages. Le Métyayer’s (1988) ACE system is a two-stage system that first converts FP programs (Backus 1978) to recursive FP programs describing the number of recursive calls of the source program, then attempts to transform the result using various program-transformation techniques to obtain a closed form. Shultis (1985) defines a denotational semantics for a simple higher-order language that models both the value and the cost of an expression. As a part of the cost model, he develops a system of “tolls,” which play a role similar to the potentials we define in our work. The tolls and the semantics are not used directly in calculations, but rather as components in a logic for reasoning about them. Sands (1990) puts forward a translation scheme in which programs in a source language are translated into programs in the same language that incorporate cost information; several source languages are discussed, including a higher-order call-by-value language. Each identifier f in the source language is associated to a *cost closure* that incorporates information about the value f takes on its arguments; the cost of applying f to arguments; and arity. Cost closures are intended to address the same issue our higher-type potentials do: recording information about the future cost of a partially-applied function. Van Stone (2003) annotates the operational semantics for a higher-order language with cost information. She then defines a category-theoretic denotational semantics that uses “cost structures” to capture cost information and shows that the latter is sound with respect to the former. Benzinger (2004) annotates NuPRL’s call-by-name operational semantics with complexity estimates. The language for the annotations is left somewhat open so as to allow greater flexibility. The analysis of the costs is then completed using a combination of NuPRL’s proof generation and Mathematica. In all of these approaches the cost domain incorporates information about values in the source language so as to provide exact costs. Our approach provides a uniform framework that can be more or less precise about the source language values that are represented. While we can implement a version that handles exact costs, we can also implement a version in which we focus just on upper bounds, which we might hope leads to simpler recurrences.

8. Conclusions and Further Work

We have described a denotational complexity analysis for a higher-order language with a general form of inductive datatypes that yields an upper bound on the cost of any well-typed program in terms of the size of the input. The two steps are to translate each source-language program e into a program $\llbracket e \rrbracket$ in a complexity language, which makes costs explicit, and then to abstract values to sizes. A consequence of the bounding theorem is that the cost component of $\llbracket e \rrbracket$ is an upper bound on the evaluation cost of e . The bounding theorem is purely syntactic and therefore applies in all models of the complexity language. By varying the semantics of the complexity language (and in particular, the notion of size), we can perform analyses at different levels of granularity. We give several different choices for the notion of size, but ultimately this is too important a decision to take out of the hands of the user through automation.

The complexity translation of Section 3 can easily be adapted to other cost models. For example, we could charge different amounts for different steps. Or, we could analyze the work and span of parallel programs by taking \mathbf{C} to be series-parallel cost graphs, something we plan to investigate in future work.

Another direction for future work is to handle different evaluation strategies. Compositionality is a thorny issue when considering call-

by-need evaluation and lazy datatypes, and as noted by Okasaki (1998), it may be that amortized cost is at least as interesting as worst-case cost. Sands (1990), Van Stone (2003), and Danielsson (2003) address laziness in their work, and as we already noted, RAML already performs amortized analyses.

We plan to extend the source language to handle general recursion. Part of the difficulty here is that the bounding relation presupposes termination of the source program (so that the derivation of $e \Downarrow^n v$, and hence cost, is well-defined). One approach would be to require the user to supply a termination proof. Or, one could define the operational semantics of the source language co-inductively (as done by, e.g., (Leroy and Grall 2009)), thereby allowing explicitly for non-terminating computations. Another approach is to adapt the partial big-step operational semantics described by Hoffmann et al. (2012). Since our source language supports inductive datatype definitions of the form `datatype strm = Cons of unit → nat × strm`, adding general recursion will force us to understand how our complexity semantics plays out in the presence of what are essentially coinductively defined values. One could also hope to prove termination in the source language by first extracting complexity bounds and then proving that these bounds in fact define total functions. Another interesting idea along these lines would be to define a complexity semantics in which the cost domain is two-valued, with one value representing termination and the other non-termination (or maybe more accurately, known termination and not-known-termination); such an approach might be akin to an abstract interpretation based approach for termination analysis.

The programs $\|e\|$ are complex higher-order recurrences that call out for solution techniques. Benzinger (2004) addresses this idea, as do Albert et al. (2011, 2013) of the COSTA project. Another relevant aspect of the COSTA work is that their cost relations use non-determinism; it would be very interesting to see if we could employ a similar approach instead of the maximization operators that we used in our examples. Ultimately we should have a library of tactics for transforming the recurrences produced by the translation function to closed (possibly asymptotic) forms when possible.

References

- E. Albert, P. Arenas, S. Genaim, and G. Puebla. Closed-form upper bounds in static cost analysis. *Journal of Automated Reasoning*, 46:161–203, 2011. doi: 10.1007/s10817-010-9174-1.
- E. Albert, P. Arenas, S. Genaim, G. Puebla, and D. Zanardini. Cost analysis of object-oriented bytecode programs. *Theoretical Computer Science*, 413(1):142–159, 2012. doi: 10.1016/j.tcs.2011.07.009.
- E. Albert, S. Genaim, and A. N. Masud. On the inference of resource usage upper and lower bounds. *ACM Transactions on Computational Logic*, 14(3):22:1–22:35, 2013. doi: 10.1145/2499937.2499943.
- J. Backus. Can programming be liberated from the von Neumann style? A functional style and its algebra of programs. *Communications of the Association for Computing Machinery*, 21(8):613–641, 1978. doi: 10.1145/359576.359579.
- R. Benzinger. Automated higher-order complexity analysis. *Theoretical Computer Science*, 318(1-2):79 – 103, 2004. doi: 10.1016/j.tcs.2003.10.022.
- N. A. Danielsson. Lightweight semiformal time complexity analysis for purely functional data structures. In A. Aiken and G. Morrisett, editors, *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 133–144. ACM Press, 2003. doi: 10.1145/1328438.1328457.
- N. Danner and J. S. Royer. Two algorithms in search of a type system. *Theory of Computing Systems*, 45(4):787–821, 2009. doi: 10.1007/s00224-009-9181-y.
- N. Danner, J. Paykin, and J. S. Royer. A static cost analysis for a higher-order language. In M. Might and D. V. Horn, editors, *Proceedings of the 7th workshop on Programming languages meets program verification*, pages 25–34. ACM Press, 2013. doi: 10.1145/2428116.2428123.
- N. Danner, D. R. Licata, and R. Ramyaa. Denotational cost semantics for functional languages with inductive types. arXiv:XXXX.XXXX, 2015.
- S. K. Debray and N.-W. Lin. Cost analysis of logic programs. *ACM Transactions on Programming Languages and Systems*, 15(5):826–875, 1993. doi: 10.1145/161468.161472.
- R. Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, 2013.
- J. Hoffmann and M. Hofmann. Amortized resource analysis with polynomial potential: A static inference of polynomial bounds for functional programs. In A. D. Gordon, editor, *Programming Languages and Systems: 19th European Symposium on Programming, ESOP 2010*, volume 6012 of *Lecture Notes in Computer Science*, pages 287–306. Springer-Verlag, 2010. doi: 10.1007/978-3-642-11957-6_16.
- J. Hoffmann and Z. Shao. Automatic static cost analysis for parallel programs. In J. Vitek, editor, *Programming Languages and Systems: 24th European Symposium on Programming, ESOP 2015*, volume 9032 of *Lecture Notes in Computer Science*, pages 132–157. Springer-Verlag, 2015. doi: 10.1007/978-3-662-46669-8_6.
- J. Hoffmann, K. Aehlig, and M. Hofmann. Multivariate amortized resource analysis. *ACM Transactions on Programming Languages and Systems*, 34(3):14:1–14:62, 2012. doi: 10.1145/2362389.2362393.
- S. Jost, K. Hammond, H.-W. Loidl, and M. Hofmann. Static determination of quantitative resource usage for higher-order programs. In M. Hermenegildo, editor, *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 223–236. ACM Press, 2010. doi: 10.1145/1706299.1706327.
- D. Le Métayer. ACE: an automatic complexity evaluator. *ACM Transactions on Programming Languages and Systems*, 10(2):248–266, 1988. doi: 10.1145/42190.42347.
- X. Leroy and H. Grall. Coinductive big-step operational semantics. *Information and Computation*, 207(2):284–304, 2009. doi: 10.1016/j.ic.2007.12.004.
- E. Moggi. Notions of computation and monads. *Information And Computation*, 93(1):55–92, 1991. doi: 10.1016/0890-5401(91)90052-4.
- J. Navas, E. Mera, P. López-Garcia, and M. V. Hermenegildo. User-definable resource bounds analysis for logic programs. In V. Dahl and I. Niemelä, editors, *Proceedings of Logic Programming: 23rd International Conference, ICLP 2007*, volume 4670 of *Lecture Notes in Computer Science*, pages 348–363, 2007. doi: 10.1007/978-3-540-74610-2_24.
- C. Okasaki. *Purely Functional Data Structures*. Cambridge University Press, 1998.
- M. Rosendahl. Automatic complexity analysis. In J. E. Stoy, editor, *Proceedings of the Fourth International Conference on Functional Programming Languages and Computer Architecture*, pages 144–156. ACM Press, 1989. doi: 10.1145/99370.99381.
- D. Sands. *Calculi for Time Analysis of Functional Programs*. PhD thesis, University of London, 1990.
- J. Shultis. On the complexity of higher-order programs. Technical Report CU-CS-288-85, University of Colorado at Boulder, 1985.
- K. Van Stone. *A Denotational Approach to Measuring Complexity in Functional Programs*. PhD thesis, School of Computer Science, Carnegie Mellon University, 2003.
- P. Wadler. Views: A way for pattern matching to cohabit with data abstraction. In *Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 307–313, 1987. doi: 10.1145/41625.41653.
- P. Wadler. The essence of functional programming. In R. Sethi, editor, *Proceedings of the 19th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 1–14. ACM Press, 1992. doi: 10.1145/143165.143169.
- B. Wegbreit. Mechanical program analysis. *Communications of the Association for Computing Machinery*, 18(9):528–539, 1975. doi: 10.1145/361002.361016.