

Complexity Analysis for a Lazy Higher-Order Language

David Sands *

Department of Computing, Imperial College
180 Queens Gate, London SW7 2BZ
email: ds@uk.ac.ic.doc

Abstract

This paper is concerned with the time-analysis of functional programs. Techniques which enable us to reason formally about a program's execution costs have had relatively little attention in the study of functional programming. We concentrate here on the construction of equations which compute the time-complexity of expressions in a lazy higher-order language.

The problem with higher-order functions is that complexity is dependent on the cost of applying functional parameters. Structures called *cost-closures* are introduced to allow us to model both functional parameters *and* the cost of their application.

The problem with laziness is that complexity is dependent on *context*. Projections are used to characterise the context in which an expression is evaluated, and cost-equations are parameterised by this context-description to give a compositional time-analysis. Using this form of context information we introduce two types of time-equation: *sufficient-time* equations and *necessary-time* equations, which together provide bounds on the exact time-complexity.

1 Introduction

This paper is concerned with the time-analysis of functional programs. Techniques which enable us to reason formally about a program's execution costs have had relatively little attention in the study of functional programming. There has been some interest in the mechanisation of program cost analysis, perhaps the main examples being [Weg75, LeM85, Ros89]. These works describe systems which analyse cost by first constructing (recursive) equations which describe the time-complexity of a functional program in a strict first-order language. A closed form expression for cost is obtained in some cases by mechanised manipulation (transformation) of these equations. The average-case solution of such equations is considered in [HC88, Fla85]. We concentrate here on the *first* part of this process—the construction of equations which compute the time-complexity of a given program. For programs written a first-order strict (i.e. call-by-value) language this is very straightforward. In the first part of this paper we show how to deal with a strict higher-order language (a fuller development can be found in [San88]). In the remainder of the paper we adapt these ideas to a lazy language. This extension is based on Wadler's use of *context-analysis* in the construction of time equations for a lazy first-order language [Wad88].

The aim is to develop a calculus that enables us to reason about time-complexity. Given a program (which we will consider to be any expression, plus a set of mutually recursive function definitions), the problem is to find a means of constructing equations which describe the cost (in terms of the number of certain elementary operations) of evaluating any expression. In this paper we choose to express cost in terms of the number of non-primitive function applications. One advantage of deriving cost-equations which are themselves expressed in a functional language is that they are amenable to a rich class of program transformation and analysis techniques c.f. [LeM85, Ros89]—this paper retains the functional flavour of these approaches.

*This work was partially supported by ESPRIT Basic Research Action P3124

The paper is organised as follows. In section 2 we consider the analysis of first and higher-order strict languages. Section 3 introduces a description of *context* that will be used in the analysis of lazy languages. Section 4 presents *sufficient-time* analysis, an upper-bound analysis for a lazy first-order language, which uses contexts that describe information that is *sufficient* to compute a value. Section 5 presents *necessary-time* analysis, a corresponding lower-bound analysis. Section 6 extends these ideas to a higher-order language.

2 Strict Time Analysis

In this section we consider the analysis of strict languages. A full presentation is given in [San88].

2.1 A First Order Language

Firstly we define a simple first-order functional language. We consider a set of mutually recursive function definitions of the form $f_i(x_1, \dots, x_{n_i}) = e_i$ and an expression to be evaluated in the context of these definitions. Expressions have the following syntax:

$$e ::= f(e_1, \dots, e_j) \mid \text{ident} \mid \text{const} \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3$$

Where f is one of the user-defined functions f_i , or a strict primitive function or constructor p .

For each equation of the form $f_i(x_1, \dots, x_{n_i}) = e_i$ it is straightforward to construct an equation taking the same arguments as the original function, which computes the cost (in terms of the number of non-primitive function calls) of applying f_i to a tuple of values. The *cost equation* (or *cost-function*) is defined as: $cf_i(x_1, \dots, x_{n_i}) = 1 + \mathcal{T}[e_i]$ where \mathcal{T} is a syntax-directed abstraction given in figure 1. These rules clearly reflect the call-by-value evaluation order. For example, in the rule for application,

$$\begin{aligned} \mathcal{T}[\text{const}] &= \mathcal{T}[\text{ident}] = 0 \\ \mathcal{T}[\text{if } e_1 \text{ then } e_2 \text{ else } e_3] &= \mathcal{T}[e_1] + \text{if } e_1 \text{ then } \mathcal{T}[e_2] \text{ else } \mathcal{T}[e_3] \\ \mathcal{T}[p(e_1 \dots e_n)] &= \mathcal{T}[e_1] + \dots + \mathcal{T}[e_n] \\ \mathcal{T}[f_i(e_1 \dots e_n)] &= cf_i(e_1 \dots e_n) + \mathcal{T}[e_1] + \dots + \mathcal{T}[e_n] \end{aligned}$$

Figure 1: First-Order Strict Cost Definition

we sum the cost of evaluating the arguments, in addition to the function application. (*N.B.* We will use infix notation to ease presentation throughout this paper)

Syntax directed derivations of this form, for similar first order languages can be found in [Weg75, LeM85, Ros89]. These works focus on some automatic techniques by which the recursive cost-equations can be manipulated to achieve non-recursive equations.

Example

As a simple example of the above scheme, consider the list-append function defined as:

$$\text{append}(x, y) = \text{if null}(x) \text{ then } y \text{ else cons}(\text{hd}(x), \text{append}(\text{tl}(x), y))$$

From this definition, applying \mathcal{T} we obtain the cost-function which computes the number of non-primitive function applications:

$$\begin{aligned} \text{cappend}(x, y) &= 1 + \text{if null}(x) \text{ then } 0 \text{ else cappend}(\text{tl}(x), y) \\ &= 1 + \text{length}(x) \end{aligned}$$

The aim of the systems described in the papers cited above is to derive just such a closed-form expression, by means of program transformation. This paper focuses on the process of obtaining the

initial cost-functions, for languages using higher-order functions and laziness—a necessary precursor to the derivation of closed-form equations describing, for example, average-case complexity.

2.2 A Higher-Order Curried Language

In this section we outline a means of deriving cost programs for a higher-order language. The time-equations are derived via two mappings. The first modifies the original equations so that functional values are augmented with information needed to describe the cost of their application. The second constructs the time-equations using these modified equations.

Firstly we define our language. We have function definitions of the form $f_i e_1 \dots e_{n_i} = exp_i$ along with curried primitive functions p_i (of arity m_i). Expressions have the following syntax:

$$exp ::= exp\ e \mid e \quad e ::= \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid (exp) \mid f_i \mid p_i \mid \text{ident} \mid \text{const}$$

For each definition $f_i x_1 \dots x_{n_i} = exp_i$ we wish to construct a cost function $cf_i x_1 \dots x_{n_i} = exp'_i$ which computes the cost of applying f_i to n_i values.

Suppose we wish to construct a cost-function for an apply function defined as: $\text{apply } f\ x = f\ x$. The cost function associated with apply should have the form:

$$C_{\text{apply}}(f, x) = 1 + \text{the cost of applying } f \text{ to } x.$$

But how do we *syntactically* refer to the cost function associated with f ?

Cost-Closures

In order to reason about the cost of application of functions, as well as the functions themselves, we introduce structures called *cost-closures*. A cost-closure is a triple (f, cf, a) of a function f , its associated cost-function cf and some arity information a . Together with cost-closures we define two (left associative) infix functions \mathbb{C} and $\mathbf{c}\mathbb{C}$ which define the application of cost-closures and the cost of application. Functions \mathbb{C} and $\mathbf{c}\mathbb{C}$ satisfy:

$$(f, cf, a)\ \mathbb{C}\ e = \begin{cases} f\ e & \text{if } a = 1 \\ (f\ e, cf\ e, a - 1) & \text{otherwise} \end{cases} \quad (f, cf, n)\ \mathbf{c}\mathbb{C}\ e = \begin{cases} cf\ e & \text{if } n = 1 \\ 0 & \text{otherwise} \end{cases}$$

The arity component of the cost-closure, and its use in the definition of $\mathbf{c}\mathbb{C}$ is explained by the fact that for reasons of efficiency and simplicity, there is no evaluation of the body of a function until the function is supplied with at least the number of arguments in its definition (this avoids the potentially expensive resolution of name clashes, and is thus a feature of most functional language implementations).

Cost-closures are used in the following way. We define two syntax-directed translation functions \mathcal{V} and \mathcal{T} . The purpose of \mathcal{V} (figure 3) is to modify the original program so that all functional objects are translated into cost-closures, and to perform application via \mathbb{C} . \mathcal{T} (figure 4) defines the cost-functions, using $\mathbf{c}\mathbb{C}$. The cost of evaluating any expression exp with respect to definitions $f_i x_1 \dots x_{n_i} = exp_i$, $i = 1, \dots, k$ is then defined by the program given in figure 2.

```

let  f'_1 x_1 ... x_{n_1} = \mathcal{V}[e_1]
      \vdots
      cf_1 x_1 ... x_{n_1} = 1 + \mathcal{T} \circ \mathcal{V}[e_1]
      \vdots
in   \mathcal{T} \circ \mathcal{V}[e]

```

Figure 2: Higher-Order Cost-Program Scheme

$$\begin{aligned}
\mathcal{V}[exp\ e] &= \mathcal{V}[exp] \circledast \mathcal{V}[e] \\
\mathcal{V}[\text{if } e_1 \text{ then } e_2 \text{ else } e_3] &= \text{if } \mathcal{V}[e_1] \text{ then } \mathcal{V}[e_2] \text{ else } \mathcal{V}[e_3] \\
\mathcal{V}[(exp)] &= (\mathcal{V}[exp]) \\
\mathcal{V}[f_i] &= (f'_i, cf_i, n_i) \quad \mathcal{V}[p_i] = (p_i, cp_i, m_i) \\
\mathcal{V}[const] &= const \quad \mathcal{V}[ident] = ident
\end{aligned}$$

Figure 3: Function Modification Map, \mathcal{V}

$$\begin{aligned}
\mathcal{T}[exp' \circledast e'] &= \mathcal{T}[exp'] + \mathcal{T}[e'] + (exp' \text{ c}\circledast e') \\
\mathcal{T}[\text{if } e'_1 \text{ then } e'_2 \text{ else } e'_3] &= \mathcal{T}[e'_1] + \text{if } e'_1 \text{ then } \mathcal{T}[e'_2] \text{ else } \mathcal{T}[e'_3] \\
\mathcal{T}[(exp')] &= (\mathcal{T}[exp']) \\
\mathcal{T}[(p_i, cp_i, m_i)] &= \mathcal{T}[(f_i, cf_i, n_i)] = \mathcal{T}[const] = 0
\end{aligned}$$

Figure 4: Cost-Expression Construction Map, \mathcal{T}

\mathcal{V} is defined on the structure of expressions exp and e . \mathcal{T} is consequently defined over the syntax of expressions generated by \mathcal{V} .

Some Optimisations

The code derived by the above translation schemes is rather more cumbersome than is necessary. This is because we introduce more \circledast 's and $\text{c}\circledast$'s than are necessary. Some straightforward optimisations simplify the cost program considerably, and can be defined according to the syntactic structure of expressions [San88].

Example

The following simple example illustrates the derivation (and the optimisation):

```
map f x = if (null x) then nil else (cons (f (hd x)) (map f (tl x)))
```

The cost-function derived from this is:

```
cmap f x = 1 + ((null,cnull,1) c\circledast x) +
  if ((null,cnull,1) \circledast x) then nil
  else (((cons,ccons,2) c\circledast (f \circledast ((hd,chd,1) \circledast x))) +
((cons,ccons,2)\circledast(f \circledast ((hd,chd,1)\circledast x)) c\circledast ((map',cmap,2)\circledast f \circledast ((tl,ctl,1)\circledast x)))
+ f c\circledast ((hd,chd,1)\circledast x) + ((hd,chd,1)c\circledast x) + ((map',cmap,2)\circledast f c\circledast ((tl,ctl,1)\circledast x)
+ ((map',cmap,2)c\circledast f) + ((tl,ctl,1)c\circledast x) )
```

Using simple optimisation schemes, we get the equivalent cost-function definition:

```
cmap f x = 1 + if (null x) then 0 else (f c\circledast (hd x)) + (cmap f (tl x))
```

2.3 Correctness

The derived program computes the number of times a certain “step” is performed in the evaluation of the program. In [San88] we formalise our intuitive model of “evaluation steps” via an *operational semantics* and prove that the number of steps our derived program computes is correct with respect to the actual operational behaviour of the original program.

3 Lazy Time Analysis: Describing Context

A major obstacle in the time-analysis of lazy languages is the problem of *context sensitivity*: the cost of evaluating an expression depends on the context in which it is used. In order to give a *compositional* treatment of the analysis of lazy-evaluation we must take into account some description of context.

3.1 Modelling Contexts with Projections

The formulation of a context which will be used in our time analysis is that provided by Wadler and Hughes [WH87] in the analysis of *strictness*. Wadler shows how this formulation of context can be useful for time analysis in [Wad88]. Here we provide an introduction to the use of *projections* to model contexts. For a fuller development the reader is referred to [WH87]; a more formal development is given in [DW89].

The basic problem is, given a function, how much information do we require from the argument in order to determine a certain amount of information about the result. Projections, in the domain theoretic sense, can provide a concise description of both the amount of information which is *sufficient* and the amount which is *necessary*.

DEFINITION 3.1 *A projection, α , is a continuous function from a domain \mathcal{D} onto itself, such that $\alpha \sqsubseteq \text{ID}_{\mathcal{D}}$ and $\alpha \circ \alpha = \alpha$, where $\text{ID}_{\mathcal{D}}$ is the identity function on \mathcal{D}*

In other words, given an object u , a projection removes information from that object ($\alpha u \sqsubseteq u$), but once this information has been removed further application has no effect ($\alpha(\alpha u) = \alpha u$). A projection is used to represent a context, where the information removed represents information not needed by that context.

In the following the terms *projection* and *context* will be synonymous, and will be ranged over by α and β .

DEFINITION 3.2 Safe Projections: *Given a (first order) function, f , of n arguments, if*

$$\alpha(f(u_1, \dots, u_n)) = \alpha(f(u_1, \dots, (\beta u_i), \dots u_n))$$

for all objects u_1, \dots, u_n , then we say that in context α , β is a safe context for the i 'th argument of f . This is abbreviated by $f^i : \alpha \Rightarrow \beta$.

Lifted Projections

We will require that projections describe two types of information: what information is *sufficient*, and what information is *necessary*. In order to describe the latter, Wadler and Hughes introduce a new domain element, \perp_{γ} , called "abort". The interpretation of $\alpha u = \perp_{\gamma}$ is that context α requires a value more defined than u . To make this work, we must have $\perp_{\gamma} \sqsubseteq \perp$ and all functions are naturally-extended to be strict in \perp_{γ} , i.e., $f(u_1, \dots, \perp_{\gamma}, \dots, u_n) = \perp_{\gamma}$. These technical devices are explained more formally in [Bur90] in terms of *lifting*.

The Projection Lattice

A projection $\alpha : \mathcal{D}_{\perp_{\gamma}} \rightarrow \mathcal{D}_{\perp_{\gamma}}$, is called a *projection over \mathcal{D}* . Projections over any domain form a lattice, with ordering \sqsubseteq , containing at least the following points:

ID	•	ID u =	u
ABS •	•	STR u =	$\begin{cases} \perp_{\gamma} & \text{if } u = \perp \text{ or } u = \perp_{\gamma} \\ u & \text{otherwise} \end{cases}$
	•	ABS u =	$\begin{cases} \perp_{\gamma} & \text{if } u = \perp_{\gamma} \\ \perp & \text{otherwise} \end{cases}$
	•	FAIL u =	\perp_{γ}

DEFINITION 3.3 A strict projection is any projection α such that $\alpha(\perp) = \perp$

The largest of such projections is STR, giving us an alternative definition of strict projections: a projection α is strict if and only if $\alpha \sqsubseteq \text{STR}$. Of the non-strict projections, the smallest is the projection ABS. This context is important since if it is safe to evaluate an expression in the context ABS, then the value of the expression will not be needed. FAIL is the unsatisfiable context.

There may be infinitely many projections, which are all either strict ($\text{FAIL} \sqsubseteq \alpha \sqsubseteq \text{STR}$) or non-strict ($\text{ABS} \sqsubseteq \alpha \sqsubseteq \text{ID}$). The four projections above will be used “polymorphically” to represent the corresponding projection over the appropriate domain.

Given two projections $\alpha \sqsubseteq \beta$, α represents a more precise description of a context than β . The context ID is therefore the least informative. Furthermore, if it is safe to evaluate some expression in a context α , then it is always safe to evaluate in a context β , $\alpha \sqsubseteq \beta$.

Contexts for Lists

The following projections are useful for building contexts over the non-flat domain of lists \mathcal{D}^* of elements from some domain \mathcal{D} :

$$\text{NIL } u = \begin{cases} \text{nil} & \text{if } u = \text{nil} \\ \perp & \text{otherwise} \end{cases} \quad \text{CONS } \alpha \beta u = \begin{cases} \text{cons}(\alpha x)(\beta xs) & \text{if } u = \text{cons } x \text{ } xs \\ \perp & \text{otherwise} \end{cases}$$

NIL is the context which requires an empty-list, and $\text{CONS } \alpha \beta$ is the context which requires a non-empty list whose head is needed in context α and whose tail is needed in context β . For example, the context (CONS STR ABS) requires a non-empty list whose first element is needed, and the rest is not.

Context Analysis

Analysing context is a *backwards analysis* [AH87]. Given a context α for a function f , what can we say about the contexts of the arguments? We need to propagate the information about the result of a function *backwards* to its arguments. i.e., given a function f of arity n , and a context α we need to find each β_i such that $f^i : \alpha \Rightarrow \beta$.

Ideally we need to find the smallest β_i , since these describe the contexts most precisely. In order to give a computable approximation we may settle for *some* β_i satisfying the above property.

Projection Transformers

A function of α yielding such a β_i is called a *projection transformer*. We will adopt the following notation: The projection-transformer written $f^{\#i}$ is a function satisfying $f^i : \alpha \Rightarrow f^{\#i} \alpha$. *N.B.* Strictly speaking we should distinguish between the syntactic objects—the program defining f , and the semantic objects—the projections, and the denotations given by some semantic function. Following the style of [WH87] we will mix these entities for notational convenience.

Rules for defining recursive equations for the projection transformers are given in [WH87]—an important result here is that a solution to these equations can be determined *automatically* if we work with finite lattices of projections, although it is not difficult to modify the equations to give more accurate projection equations (which are harder to solve).

4 Sufficient-Time Analysis

In this section we show how context information can be used to aid the time analysis of a lazy first-order language; Sufficient-time analysis (with some minor differences) corresponds to the time analysis presented in [Wad88]. The information obtained by the backwards analysis is used to derive equations which compute an upper bound to the precise cost of a given program. This upper-bound

$$\begin{aligned}
\mathcal{T}_s[\text{const}]\alpha &= \mathcal{T}_s[\text{ident}]\alpha = 0 \\
\mathcal{T}_s[\text{if } e_1 \text{ then } e_2 \text{ else } e_3]\alpha &= \alpha \hookrightarrow_s \mathcal{T}_s[e_1]\text{ID} + \text{if } e_1 \text{ then } \mathcal{T}_s[e_2]\alpha \text{ else } \mathcal{T}_s[e_3]\alpha \\
\mathcal{T}_s[p(e_1 \dots e_n)]\alpha &= \mathcal{T}_s[e_1](p^{\#1}\alpha) + \dots + \mathcal{T}_s[e_n](p^{\#n}\alpha) \\
\mathcal{T}_s[f_i(e_1 \dots e_n)]\alpha &= cf_i(e_1 \dots e_n, \alpha) + \mathcal{T}_s[e_1](f_i^{\#1}\alpha) + \dots + \mathcal{T}_s[e_n](f_i^{\#n}\alpha)
\end{aligned}$$

Figure 5: Definition of \mathcal{T}_s

is obtained by using information which tells us what values are *sufficient* to compute an expression. We call the resulting analysis a *sufficient-time* analysis.

4.1 Context-Parameterised Cost Functions

As in the first-order time analysis of section 2, we will define a *cost-function*, cf_i , for each function f_i defined in the original program. As before the cost functions will take as parameters the original arguments to the functions, but in addition they will be parameterised by a *context*, representing the context in which the functions are evaluated.

How can cost-functions make use of context ?

We know that any expression in the context ABS will be ignored, so the cost in this context is zero. In any other context the cost of a function application will be (approximated above by) 1 + the cost of evaluating the body of the function, in that context.

We define the cost functions associated with each function $f_i(x_1 \dots x_{n_i}) = e_i$ to be

$$cf_i(x_1, \dots, x_{n_i}, \alpha) = \alpha \hookrightarrow_s 1 + \mathcal{T}_s[e_i]\alpha$$

where we introduce the notation $\alpha \hookrightarrow_s e$ to abbreviate cost e “guarded” by context α :

$$\alpha \hookrightarrow_s e = \begin{cases} 0 & \text{if } \alpha = \text{ABS} \\ e & \text{otherwise} \end{cases}$$

The syntactic map \mathcal{T}_s defined in figure 5 is very similar to that defined in figure 1, but is defined with respect to a particular context. $\mathcal{T}_s[e]\alpha$ defines the cost of evaluating expression e in context α . It makes use of the context transformers $f_i^{\#1} \dots f_i^{\#n_i}$ defined for each function f_i , which satisfy the required safety criterion. In particular it will be appropriate to set $f^{\#i}(\text{ABS}) = \text{ABS}$, since if the result of a function is not needed, then neither are its arguments.

The rule for function application tells us that the cost of evaluating a function application is the associated cost-function applied to the arguments (and the context) plus the sum of evaluating the arguments in the contexts prescribed by the context-transformers.

The conditional expression, like any other, has zero cost in the context ABS (guaranteed by the use of \hookrightarrow_s). Otherwise we sum the cost of evaluating the condition (which may or may not be evaluated, hence the safe-context for boolean values ID, c.f. [Wad88]) plus either the cost of the alternate or the consequent, depending on the *value* of the condition.

The cost of evaluating any expression in the context ABS is zero, so we have:

PROPOSITION 4.1 For every expression e , $\mathcal{T}_s[e]\text{ABS} = 0$

PROOF Straightforward structural induction in e □

A Small Example

Consider the program: $\text{hd}(\text{cons}(\text{not}(\text{true}), \text{exp}))$, where $\text{not}(x) = \text{if } x \text{ then false else true}$ and exp represents some arbitrary expression. The cost-function for not is

$$\begin{aligned} \text{cnot}(x, \alpha) &= \alpha \hookrightarrow, 1 + (\alpha \hookrightarrow, 0 + \text{if } x \text{ then } 0 \text{ else } 0) \\ &= \alpha \hookrightarrow, 1 \end{aligned}$$

We assume a boolean-valued program is evaluated in the context STR , and so the cost program is defined by: $\mathcal{T}_s[\text{hd}(\text{cons}(\text{not}(\text{true}), \text{exp}))] \text{STR}$, which is, by definition

$$\text{cnot}(\text{true}, \text{cons}^{\#1}(\text{hd}^{\#1}(\text{STR}))) + 0 + \mathcal{T}_s[\text{exp}] \text{cons}^{\#2}(\text{hd}^{\#1}(\text{STR}))$$

The context transformers for the primitive functions satisfy

$$\text{hd}^{\#1}(\alpha) = \text{CONS } \alpha \text{ ABS} \quad \text{cons}^{\#1}(\text{CONS } \alpha \beta) = \alpha \quad \text{cons}^{\#2}(\text{CONS } \alpha \beta) = \beta$$

and so the cost is $\text{cnot}(\text{true}, \text{STR}) + \mathcal{T}_s[\text{exp}] \text{ABS} = 1$, for any expression exp .

4.2 Approximation and Safety

What are the precise properties of the cost programs? Here we consider the approximation and correctness properties of the “lazy” cost-program.

Approximation

The expression $\mathcal{T}_s[e]\alpha$ gives an upper-bound estimate to the cost of lazy evaluation of e in context α . The cost expressions formed by \mathcal{T}_s are a refinement of the call-by-value cost-expression (section 2) in which subexpressions whose values are not needed do not contribute to the cost equation. Since the safety condition for projections does not specify that we require the smallest possible projection, the context ABS may be approximated by any larger projection. This approximation is reflected in the cost-program as an over-estimation of cost. (In the extreme case the context transformers are such that the context ABS is *never* derived in the cost-program, and so the value of the cost program is the same as that given by the strict derivation of figure 1.) Note also that in computing the cost of a function application $f(e)$ in context α the cost due to e will only be counted *once*. The context of e , $f^{\#1}(\alpha)$ will be the net context of the possible contexts in which e is shared, and so the process properly models call-by-need.

Safety

Whenever the cost-program terminates yielding a value, that value is indeed an upper bound to the time cost of evaluating the program lazily. A problem with this analysis method is that there are cases when the cost-program does not yield a value when it should do so. Firstly the cost-program may not terminate even when the program does—non-terminating cost expressions can be thought of as “computing” the worst possible upper-bound to the cost. However the approximation in the cost-program can lead to arbitrary run-time errors (i.e. not just nontermination). In the next section we introduce *necessary-time* equations which allow us to place a lower-bound on the precise complexity and which have better termination properties.

5 Necessary-Time Analysis

So far we have outlined the use of contexts to derive equations which can give an upper-bound to the time-complexity of an expression in a particular context. As mentioned previously, this idea is based on [Wad88]. The cost-functions which compute this sufficient-complexity are only partially correct in the sense that if they compute a value, then that value is indeed an upper-bound to the time-cost of a program. There is potentially much more information about context using the projections described:

$$\begin{aligned}
\mathcal{T}_n[\text{const}]\alpha &= \mathcal{T}_n[\text{id}]\alpha = 0 \\
\mathcal{T}_n[\text{if } e_1 \text{ then } e_2 \text{ else } e_3]\alpha &= \alpha \hookrightarrow_n \mathcal{T}_n[e_1]\text{STR} + \text{if } e_1 \text{ then } \mathcal{T}_n[e_2]\alpha \text{ else } \mathcal{T}_n[e_3]\alpha \\
\mathcal{T}_n[p(e_1 \dots e_n)]\alpha &= \mathcal{T}_n[e_1](p^{\#1}\alpha) + \dots + \mathcal{T}_n[e_n](p^{\#n}\alpha) \\
\mathcal{T}_n[f_i(e_1 \dots e_n)]\alpha &= cf_i(e_1 \dots e_n, \alpha) + \mathcal{T}_n[e_1](f_i^{\#1}\alpha) + \dots + \mathcal{T}_n[e_n](f_i^{\#n}\alpha)
\end{aligned}$$

Figure 6: Definition of \mathcal{T}_n

strict contexts allow us to describe the amount of information which is *necessary* to compute a value. In this section we show how the use of this information can give us equations which describe a lower bound to the precise time-cost (the *necessary-time*) and which overcome the termination deficiencies of sufficient-time analysis. The key to sufficient-time analysis is the use of the context ABS to deduce that an expression will not be evaluated. The key to *necessary-time* is the operational interpretation of the *strict* projections.

5.1 Necessary-Cost Functions

In order to construct functions which compute the necessary-cost of evaluating a function in a particular context, we make the following operational connection between expressions which can be safely evaluated in a strict context, and their operational behaviour.

- If it is safe to evaluate an expression of the form $f(e_1, \dots, e_n)$ in a strict context, then operationally, we know that this outermost application must be reduced.

Conversely, if an expression is evaluated in a non-strict context then that expression *may or may not* be reduced (only the context ABS allows us to conclude that it *definitely* will not).

Motivated by this observation, we now define the necessary-cost. The cost of evaluating an expression e in a context α is given by $\mathcal{T}_n[e]\alpha$ where \mathcal{T}_n is once again a mapping defined over the syntax of expressions, and assuming some safe context transformers for the user-defined functions.

For each function definition of the form $f_i(x_1 \dots x_{n_i}) = e_i$ we will define an associated necessary-cost-function $cf_i(x_1, \dots, x_{n_i}, \alpha) = \alpha \hookrightarrow_n 1 + \mathcal{T}_n[e_i]\alpha$, where we use the notation $\alpha \hookrightarrow_n e$ to abbreviate necessary-cost e modulo context α :

$$\alpha \hookrightarrow_n e = \begin{cases} e & \text{if } \alpha \sqsubseteq \text{STR} \\ 0 & \text{otherwise} \end{cases}$$

The definition of \mathcal{T}_n is given in figure 6. The rules are very similar to the definitions for \mathcal{T}_s but we use \hookrightarrow_n in place of \hookrightarrow_s . The only other difference is in the translation for the conditional expression.

PROPOSITION 5.1 *For all contexts α , if $\alpha \sqsubseteq \text{STR}$ then*

$$\alpha(\text{if } u_1 \text{ then } u_2 \text{ else } u_3) = \alpha(\text{if } \text{STR}(u_1) \text{ then } u_2 \text{ else } u_3)$$

PROOF Straightforward by cases according, $u_1 \sqsupseteq \perp$ and $u_1 \sqsubseteq \perp$ □

This tells us that in any strict context it is safe to evaluate the condition in the context STR, and thus gives us the appropriate context for determining the cost due to the condition in the conditional expression.

5.2 Example

As a example of necessary-time analysis we use insertion-sort (as in [Wad88]). The definitions are given in figure 7. The necessary-time equations constructed according to \mathcal{T}_n are given in figure 8: In this example we wish to consider the cost of evaluating min in a strict context. We are not particularly concerned here with the techniques for deriving the safe projection transformers. We

```

insert(x,xs) = if null(xs) then cons(x,nil)
              else if x < hd(xs) then cons(x,xs)
                else cons(hd(xs),insert(x,tl(xs)))
sort(xs)     = if null(xs) then nil
              else insert(hd(xs),sort(tl(xs)))
min(xs)      = hd(sort(xs))

```

Figure 7: Insertion Sort

```

cinsert(x,xs, $\alpha$ ) =  $\alpha \hookrightarrow_n 1$  + if null(xs) then 0
                    else if x < hd(xs) then 0
                      else cinsert(x,tl(xs),cons#2( $\alpha$ ))
csort(xs, $\alpha$ )     =  $\alpha \hookrightarrow_n 1$  + if null(xs) then 0
                    else cinsert(hd(xs),sort(tl(xs)), $\alpha$ ) +
                      csort(tl(xs),insert#2( $\alpha$ ))
cmin(xs, $\alpha$ )      =  $\alpha \hookrightarrow_n 1$  + csort(xs,hd#1( $\alpha$ ))

```

Figure 8: Necessary-Cost Functions

note however that the projection transformers needed in this example are members of the finite domains for lists (and integers) described in [WH87] for the purpose of strictness analysis, and as such can be determined mechanically by fixpoint iteration. The equations we require are:

```

hd#1(STR) = CONS STR ABS
cons#2(CONS STR ABS) = ABS
insert#2(CONS STR ABS) = CONS STR ABS

```

Now we examine the cost of min:

```

cmin(xs,STR) = STR  $\hookrightarrow_n 1$  + csort(xs,hd#1(STR))
              = 1 + csort(xs,CONS STR ABS)
csort(xs,CONS STR ABS) = 1 + if null(xs) then 0
                        else cinsert(hd(xs),sort(tl(xs)),CONS STR ABS) +
                          csort(tl(xs),insert#2(CONS STR ABS))
                        = 1 + if null(xs) then 0
                          else cinsert(hd(xs),sort(tl(xs)),CONS STR ABS) +
                            csort(tl(xs),CONS STR ABS)
cinsert(y,ys,CONS STR ABS) = 1 + if null(ys) then 0
                             else if y < hd(ys) then 0
                               else cinsert(y,tl(ys),cons#2(CONS STR ABS))
                             = 1 + if null(ys) then 0
                               else if y < hd(ys) then 0
                                 else cinsert(y,tl(ys),ABS)
                             = 1

```

```

and so csort(xs,CONS STR ABS) = 1 + if null(xs) then 0
                               else 1 + csort(tl(xs),CONS STR ABS)

```

This simple recurrence has the exact solution $1 + 2 \cdot \text{length}(xs)$ and so

$\text{cmin}(xs, \text{STR}) = 2 + 2 * \text{length}(xs)$

In this example the sufficient-time equations derive the same result, since the contexts (CONS STR ABS) and STR are very precise (i.e. they are the smallest safe projections). Therefore we can conclude that this is the *exact* time complexity.

5.3 Approximation and Safety

The expression $\mathcal{T}_n[e]\alpha$ gives a lower-bound estimate to the cost of the lazy evaluation of e in context α . For a non-strict context α the lower bound must be zero since an expression in such a context *may or may not* need to be evaluated. Proposition 5.2 below establishes this property.

PROPOSITION 5.2 *For every expression e , $\text{ABS} \sqsubseteq \alpha \sqsubseteq \text{ID} \Rightarrow \mathcal{T}_n[e]\alpha = 0$*

PROOF Structural induction in e . □

Safety

The necessary-cost programs enjoy better termination properties than the sufficient-cost programs, being at least as well defined as the original program. We state this property in the following way:

THEOREM 5.3 *Given mutually recursive functions f_1, \dots, f_m , defined by equations:*

$$f_i(x_1, \dots, x_{n_i}) = e_i, \quad i = 1 \dots m$$

then for all objects u_1, \dots, u_{n_i} , and contexts α

$$\alpha(f_i(u_1, \dots, u_{n_i})) \sqsupseteq \perp \Rightarrow cf_i(u_1, \dots, u_{n_i}, \alpha) \sqsupseteq \perp$$

where cf_i is defined by the equation $cf_i(x_1, \dots, x_{n_i}, \alpha) = \alpha \hookrightarrow_n \mathcal{T}_n[e_i]\alpha$

PROOF Omitted—a fixed point induction over the functions and cost-functions simultaneously. □

6 Higher-Order Lazy Time-Analysis

In this section we develop an extension to the techniques for lazy time analysis to incorporate higher-order functions. This is achieved by adaptation of the higher-order analysis given in section 2, illustrated with a conservative extension to the context information available for first-order functions.

6.1 Context Information

The extension of lazy-time analysis to higher-order functions also needs context information. Here we immediately run into some problems. The techniques which we have assumed so far, concerning the form and derivation of context transformers, cannot be directly extended to higher-order functions. Consider, for example, an instance of the apply function, $\text{apply } f \ x$, in some context α . The problem here is that there is no useful context information that can be propagated to x (by any context function $\text{apply}^{\#2}$) which is *independent* of the function f .

Wray's thesis [Wra86] shows how to handle a "second order" language (for strictness analysis) by additional parameterization of the context transformers to include the context transformers for functional arguments. An approach to fully higher-order backwards analysis is outlined in [Hug87]. This is based on a mixture of abstract interpretation (forwards analysis) and first-order backwards analysis. For the purposes of this section it will not be necessary to introduce these devices. Instead we will demonstrate our methods with a sufficient-time analysis using a very simple extension of the context information to higher-order functions. It is expected that the information provided by a full development of context analysis for higher-order functions could be accommodated in the time analysis we present here.

The language we use here is defined by the same grammar as that of the higher-order language in section 2.

6.2 The Projection Transformers

The method we shall describe for constructing the time equations will require the use of the same style of projection transformers that are used for the first-order analysis—for each function definition f_i we will require projection transformers $f_i^{\#k}$ such that $f_i^{\#k}\alpha \Rightarrow (f_i^{\#k}\alpha)$.

Since we are working with a higher-order language, we may expect expressions of the form

$$f_i e_1 \dots e_{n_i} e_{n_i+1} \dots e_m$$

Here the contexts propagated to expressions $e_1 \dots e_{n_i}$ are determined by the projection transformers of f_i . For a conservative estimate we know it is safe to propagate the context ID to the expressions $e_{n_i+1} \dots e_m$. In fact, the analysis we present will be able to use more precise information in this instance.

Objects of function type will also require projections to describe the context in which they are needed. A projection of a function gives a function which has less defined results on some of its arguments. For the purpose of time analysis it is sufficient to use the four-point context domain to describe the amount of evaluation of a functional argument (*i.e.* all or nothing). In an expression of the form $exp\ e$ in a context α , we can safely set the context for exp to be a mapping of α into the four-point domain for functions. For convenience we define a functional \diamond to perform this task:

$$\diamond\alpha = \begin{cases} \text{FAIL} & \text{if } \alpha = \text{FAIL} \\ \text{ABS} & \text{if } \alpha = \text{ABS} \\ \text{STR} & \text{if } \text{FAIL} \sqsubset \alpha \sqsubseteq \text{STR} \\ \text{ID} & \text{if } \text{ABS} \sqsubset \alpha \sqsubseteq \text{ID} \end{cases}$$

6.3 Accumulating Cost-Functions

As in the strict higher-order language we will define for each function in the language a cost function, constructed via two syntactic maps. The first, \mathcal{V}_L , plays the same rôle as that of \mathcal{V} in the higher-order strict language — it constructs cost-closures and makes their application explicit via an apply function \mathcal{O}_L . The second, \mathcal{T}_L , is used to define the cost-expressions. In the following we use the term *cost-expression* to refer to objects of type $context \rightarrow cost$. The definitions of \mathcal{V}_L and \mathcal{T}_L are given in figures 9 and 10. These definitions will be explained in the following sections.

User-defined functions

For each function defined $f_i\ x_1 \dots x_{n_i} = e_i$ we define a sufficient-cost function to be

$$cf_i\ \langle x_1, c_1 \rangle \dots \langle x_{n_i}, c_{n_i} \rangle \alpha = \alpha \hookrightarrow_s 1 + \mathcal{T}_L \circ \mathcal{V}_L[e_i]\ \alpha + c_1(f_i^{\#1}\alpha) + \dots + c_{n_i}(f_i^{\#n_i}\alpha)$$

In addition to the context-transformers, the cost functions require modified versions of functions themselves: $f_i^{\#k}\ x_1 \dots x_{n_i} = \mathcal{V}_L[e_i]$

Application and its Cost

The cost-functions defined above now have additional parameterisation in the form of cost-expressions paired with each argument. We will explain this choice by considering the cost associated with function application $exp\ e$.

In the higher-order strict language, application is first translated to $exp' \circ e'$ (were exp' is defined according to \mathcal{V}) and the cost of evaluation is $\mathcal{T}[exp'] + \mathcal{T}[e'] + exp' \circ e'$. Suppose we begin by re-using \mathcal{V} , and we attempt to define (with respect to some context β) a lazy version of \mathcal{T} , \mathcal{T}_L .

In the rule for $\mathcal{T}_L[\text{exp}' \circledast e'] \beta$ we must propagate the context β to the appropriate cost-expressions. We can map β into a four-point domain (overloading \diamond) to get a safe context for the function exp' . We do not know the appropriate context for e' , but we can always safely use the context ID and set

$$\mathcal{T}_L[\text{exp}' \circledast e'] \beta = \mathcal{T}_L[\text{exp}'] \diamond \beta + \mathcal{T}_L[e'] \text{ID} + (\text{exp}' \text{c}\circledast e') \beta$$

Two major problems make this rule unsatisfactory.

1. No useful context information is propagated to e' . The information we have available is the projection transformers, but this is not used since we do not in general know which projection transformer is appropriate.
2. If we have a partial application, for example if exp is cons (exp' is $(\text{cons}, \text{ccons}, 2)$) then e may not be evaluated at all.

We solve both of these problems by passing both the argument, *and* the cost-expression to the cost function. It is then the cost function's task to apply the appropriate context (which is determined by the projection transformers of the function) to these cost expressions—see the cost-function scheme above. We introduce new versions of \circledast and $\text{c}\circledast$ to accommodate these requirements.

Cost-closures and the apply function

For these reasons we need to define a new version of \mathcal{V} and a different version of the function \circledast . The “lazy” version of \mathcal{V} , \mathcal{V}_L is defined in figure 9. Because, in the rule for application, the cost-closure $\mathcal{V}_L[\text{exp}]$ is applied to the cost-expression $\mathcal{T}_L[e]$, we need a new version of the \circledast function which satisfies:

$$(f, cf, a) \circledast_L (e, ce) = \begin{cases} f e & \text{if } a = 1 \\ (f e, cf (e, ce), a - 1) & \text{otherwise} \end{cases}$$

Note that cost-closures retain the same *function-costfunction-arity* structure.

Defining the cost-expressions

Figure 10 also defines cost-expressions via a mapping \mathcal{T}_L . A significant difference here is that we do not make the definition with respect to a particular context. This is because we wish to pass cost-expressions (functions $\text{context} \rightarrow \text{cost}$) to the cost-functions without applying them to a particular context.

To define \mathcal{T}_L we define a couple of useful functions:

- Addition of cost expressions: we use a specialised addition operator, $\diamond+$, which (for the left operand) maps the context into the four-point projection domain of the left operand: $(ce_1 \diamond+ ce_2) \alpha = (ce_1 (\diamond \alpha)) + (ce_2 \alpha)$. By allowing \diamond to be polymorphic, $\diamond+$ is associative.
- The null cost-expression: the function $\bar{0}$ gives zero-cost in any context, so $\bar{0} \alpha = 0$ for any α .

Consider the rule for application: $\mathcal{T}_L[\text{exp}' \circledast_L (e', ce')] = \mathcal{T}_L[\text{exp}'] \diamond+ (\text{exp}' \text{c}\circledast_L (e', ce'))$

If we apply this expression to a context β , we get $\mathcal{T}_L[\text{exp}'] \diamond \beta + (\text{exp}' \text{c}\circledast_L (e', ce')) \beta$. To ensure $\text{c}\circledast_L$ gives us a cost expression, only a small change is needed in the definition of $\text{c}\circledast$

$$(f, cf, n) \text{c}\circledast_L (e, ce) = \begin{cases} cf (e, ce) & \text{if } n = 1 \\ \bar{0} & \text{otherwise} \end{cases}$$

Primitive functions

The cost-function associated with a primitive function p_i of arity m_i is

$$cp_i (x_1, c_1) \dots (x_{m_i}, c_{m_i}) \alpha = c_1 (p_i^{\#1} \alpha) + \dots + c_{m_i} (p_i^{\#m_i} \alpha)$$

$$\begin{aligned}
\mathcal{V}_L[\text{exp } e] &= \mathcal{V}_L[\text{exp}] \mathbf{c}_L (\mathcal{V}_L[e], \mathcal{T}_L \circ \mathcal{V}_L[e]) \\
\mathcal{V}_L[\text{if } e_1 \text{ then } e_2 \text{ else } e_3] &= \text{if } \mathcal{V}_L[e_1] \text{ then } \mathcal{V}_L[e_2] \text{ else } \mathcal{V}_L[e_3] \\
\mathcal{V}_L[(\text{exp})] &= (\mathcal{V}_L[\text{exp}]) \\
\mathcal{V}_L[f_i] &= (f'_i, cf_i, n_i) \quad \mathcal{V}_L[p_i] = (p_i, cp_i, m_i) \\
\mathcal{V}_L[\text{const}] &= \text{const} \quad \mathcal{V}_L[\text{ident}] = \text{ident}
\end{aligned}$$

Figure 9: The function modification map

$$\begin{aligned}
\mathcal{T}_L[\text{exp}' \mathbf{c}_L (e', ce')] &= \mathcal{T}_L[\text{exp}'] \diamond (\text{exp}' \mathbf{c}_L (e', ce')) \\
\mathcal{T}_L[\text{if } e'_1 \text{ then } e'_2 \text{ else } e'_3] &= \mathcal{T}_L[e'_1] \diamond \text{if } e'_1 \text{ then } \mathcal{T}_L[e'_2] \text{ else } \mathcal{T}_L[e'_3] \\
\mathcal{T}_L[(\text{exp}')] &= (\mathcal{T}_L[\text{exp}']) \\
\mathcal{T}_L[(p_i, cp_i, m_i)] &= \mathcal{T}_L[(f_i, cf_i, n_i)] = \mathcal{T}_L[\text{const}] = \mathcal{T}_L[\text{ident}] = \bar{0}
\end{aligned}$$

Figure 10: The cost-function construction map

Applying the above schemes in the construction of time-equations requires that we remove (partially evaluate) unnecessary instances of \mathbf{c}_L , and $\mathbf{c}\mathbf{c}_L$, as we outlined in section 2.2. In addition we need to specialise functions to remove unnecessary parameters—this is because of the additional parameterisation involved in both modified functions, and cost-functions. The (somewhat lengthy) examples have been omitted, but it is worth noting that the process could benefit from some simple mechanical support.

7 Conclusions

We have presented a method of analysing the time complexity of a lazy higher-order functional language. The techniques for a strict higher-order language are more fully developed in [San88]. We have extended of these ideas to give a treatment of lazy higher-order languages, based upon [Wad88]: projections are used to characterise the context in which an expression is evaluated, and cost-equations are parameterised on this context-description. We have introduced two types of time-equation: *sufficient-time* equations (corresponding to the equations in [Wad88]), and *necessary-time* equations, which together provide bounds on the exact time-complexity.

7.1 Related Work

A (non-compositional) means of analysing a call-by-name language is considered in [LeM88]. Le Métayer's solution involves transforming a call-by-name program into a strongly equivalent one with call-by-value semantics. The call-by-value program can be analysed using "strict" techniques (such as those presented in section 2). The translation, however, makes the program significantly more complex, and it not clear that the translation preserves the number of steps that are being counted in the analysis.

Bjerner's time analysis for programs in the language of Martin-Löf type-theory [Bje89] has relevance to the analysis of first-order lazy functional languages, and provided inspiration for Wadler's work. His operational model of contexts, *evaluation degrees* could form an alternative basis for the work presented here. More recently, Bjerner and Holmström [BH89] have adapted the ideas in [Bje89] to give a calculus for the time analysis of a first-order functional language. The equations used to describe context are *precise*, thus specifying an *exact* time-analysis. The problem here is that the equations cannot be solved mechanically. The main correctness theorem developed (independently)

in [BH89] (apart from the correctness of the context equations) corresponds very closely to theorem 5.3—if we view their model of context (called “demands”) as projections, we get a class of projections for which necessary and sufficient times will always be equal. Equations for this class of “exact” projections can be derived with a straightforward modification of the projection equations in [WH87].

7.2 Further Work

Higher-Order Context Information

The use of first-order context analysis in the analysis of a higher-order language means that, even though cost-expressions are passed as arguments so they are applied to the appropriate context, there are many cases where the contexts derived for higher-order functions are not sufficiently precise. Consider the following function definition: For satisfiable contexts α , the `apply` function (`apply f x = f x`) has the following projection transformers: `apply#1α = ◇α`, and `apply#2α = ID`.

Without knowing about the context of the function `apply`, the context for x is approximated by the least informative context `ID`.

The sufficient-time equation constructed with these projection transformers is

$$\text{apply } \langle f, fc \rangle \langle x, xc \rangle \alpha = \alpha \hookrightarrow, 1 + fc(\diamond\alpha) + xc \text{ ID} + (f \text{ c@ } \langle x, \bar{0} \rangle) \alpha$$

The lack of accurate projection transformers means that the cost-expression `xc` is applied to the imprecise context `ID`—it is not difficult to construct examples where this gives an unsatisfactory time analysis. Context-analyses for higher-order languages are not well-developed. As mentioned before, Wray’s strictness analysis handles “second order” functions—projection equations can be extended to handle such functions, and the resulting context descriptions can be used by cost-functions presented here. Fully higher-order analyses still present problems for the construction of both approximate and precise context equations.

An alternative solution to this problem further utilises the technique of “passing” cost-expressions. The expression bound to x in the function `apply` above is evaluated in the context of the function bound to f , so we can pass the cost expression on to the cost function associated with f as follows:

$$\text{apply } \langle f, fc \rangle \langle x, xc \rangle \alpha = \alpha \hookrightarrow, 1 + fc(\diamond\alpha) + (f \text{ c@ } \langle x, xc \rangle) \alpha$$

To generalise this technique we must check that any parameter whose cost-expression we wish to propagate is not shared (i.e. it is not required in more than one context). For a sufficient-time analysis we could propagate to all contexts, while in a necessary-time analysis we could choose to propagate the cost expression to a single context. In addition we need to determine when the propagation is necessary, since unnecessary propagation (i.e. when the context information is sufficiently precise) decreases the compositionality of cost-functions with no additional benefit.

Acknowledgements

Thanks to Jesper Andersen, Chris Hankin, Sebastian Hunt and Daniel Le Métayer for their useful suggestions relating to earlier drafts of this paper.

References

- [AH87] S. Abramsky and C.L. Hankin, editors. *Abstract Interpretation of Declarative Languages*. Ellis Horwood, 1987.
- [BH89] B. Bjerner and S. Holmström. A compositional approach to time analysis of first order lazy functional programs. In *Functional Programming Languages and computer architecture, conference proceedings*. ACM press, 1989.

- [Bje89] B. Bjerner. *Time Complexity of Programs in Type Theory*. PhD thesis, Chalmers University of Technology, 1989.
- [Bur90] G.L. Burn. A relationship between abstract interpretation and projection analysis (extended abstract). In *17th ACM Symposium on Principals of Programming Languages*, January 1990.
- [DW89] K. Davis and P. Wadler. Backwards strictness analysis: Proved and improved. In *Proceedings of Glasgow Workshop on Functional Programming*, August 1989.
- [Fla85] P. Flajolet. *Mathematical methods in the analysis of algorithms and data structures*. Report 400, INRIA, Le Chesnay, France, May 1985.
- [HC88] T. Hickey and J. Cohen. Automating program analysis. *J. ACM*, 35:185–220, January 1988.
- [Hug87] R. J. M. Hughes. Backwards analysis of functional programs. DoC Research Report CSC/87/R3, University of Glasgow, March 1987.
- [LeM85] D. LeMétayer. Mechanical analysis of program complexity. In *ACM SIGPLAN 85 Symposium*, July 1985.
- [LeM88] D. LeMétayer. Analysis of functional programs by program transformation. In *Second France–Japan Artificial Intelligence and Computer Science Symposium*. North-Holland, 1988.
- [Ros89] M. Rosendahl. Automatic complexity analysis. In *Functional Programming Languages and computer architecture, conference proceedings*. ACM press, 1989.
- [San88] D. Sands. Complexity analysis for a higher order language. Technical Report DOC 88/14, Imperial College, October 1988.
- [Wad88] P. Wadler. Strictness analysis aids time analysis. In *15th ACM Symposium on Principals of Programming Languages*, January 1988.
- [Weg75] B. Wegbreit. Mechanical program analysis. *C.ACM*, 18:528–539, September 1975.
- [WH87] P. Wadler and R. J. M. Hughes. Projections for strictness analysis. In *1987 Conference on Functional Programming and Computer Architecture*, Portland, Oregon, September 1987.
- [Wra86] S. C. Wray. Programming techniques for functional languages. Technical Report 92, University of Cambridge Computer Laboratory, June 1986.