

Towards a Theory of Reach

Jonathan Fowler and Graham Hutton

School of Computer Science
University of Nottingham, UK

Abstract. When testing a program, there are usually some parts that are rarely executed and hence more difficult to test. Finding inputs that guarantee that such parts are executed is an example of a *reach problem*, which in general seeks to ensure that targeted parts of a program are always executed. In previous work, Naylor and Runciman have developed a reachability solver for Haskell, based on the use of lazy narrowing from functional logic programming. Their work was focused on practical issues concerning implementation and performance. In this paper, we lay the groundwork for an underlying theory of such a system, by formally establishing the correctness of a simple reach solver.

1 Introduction

A desirable goal of software testing is for every reachable expression within a program to contribute to at least one test execution of the program. The testing then exhibits program coverage. Random property testing systems such as Quickcheck [3] often cover most of a program, but particularly hard to reach expressions may remain untested. The *Reach* system [12] was developed to address this problem, by generating inputs that execute a particular target expression. By using the Haskell Program Coverage (HPC) tool [5] to find expressions which are not tested by Quickcheck, and Reach to generate inputs that execute these expressions, the goal of program coverage can be achieved.

Work to date on the Reach system by Naylor and Runciman [12, 13] has focused on the implementation and performance of various underlying solvers. In this paper, we investigate a formal definition for the *reach problem*, and how the forward solver defined in their original paper [12] can be shown to be correct. Having such a theory is important to check the correctness of more complex solvers, such as backwards solver described in Naylor's thesis [13]. The act of formalisation also opens up new potential avenues for further research into alternate evaluation strategies, as discussed in section 10.

The forward reach solver developed by Naylor and Runciman uses a *lazy narrowing* evaluation strategy adapted from functional logic programming. Lazy narrowing can be thought of as an extension to the semantics of a non-strict language to include reduction rules for free variables. The basic idea is that when the value of a free variable is required for a case analysis to proceed, we bind the free variable to each possible alternative form that it may have. To focus on the essence of the problem, we initially consider a minimal language (section 3)

that includes only Peano-encoded natural numbers, a target expression, and case expressions. Abstracting away from the details of a real language such as Haskell we keep the presentation neat and concise but still include enough detail to express and understand the properties of the reach problem and lazy narrowing. Within the context of this minimal language we:

- Extend the language with free variables, and give a precise definition for the ‘reach problem’ in this setting (section 4);
- Define a lazy narrowing semantics for the extended language and use the semantics to define a forward reach solver (section 5);
- Show that the lazy narrowing semantics is sound and complete with respect to the original semantics, and that our reach solver is correct (section 6);
- Provide a mechanical verification of our results in the Agda system, and make the proof scripts freely available online (section 7);
- Describe how the language can be extended with a number of additional features, and extend the Agda formalisation accordingly (section 8).

We present proofs for our main results based on a number of lemmas, but for brevity do not provide proofs for the lemmas and refer the interested reader to the accompanying Agda code for the details [4]. The intended audience for the article is functional programmers with a basic knowledge of semantics. No prior knowledge of Reach is assumed; an introduction is given in section 2.

2 The Reach Problem

Reach [12] is a tool for Haskell that can be used help achieve program coverage. A *reach problem* is a Haskell program with a marked target expression and source function. The goal is to find an input to the source function that entails evaluation of the target expression. The target is typically placed in a rarely evaluated expression within the program. The inputs generated from the running of the Reach solver can then be used as test cases for these expressions.

As an example, consider a simplified version of a *balance* function from the standard library *Data.Map*. The *balance* function takes a binary tree and redistributes the tree when one sub-tree contains substantially more elements than the other, in this case four times as many:

$$\begin{aligned}
 & \textit{balance} :: \textit{Tree} \ a \ \rightarrow \ \textit{Tree} \ a \\
 & \textit{balance} \ (\textit{Leaf} \ a) = \textit{Leaf} \ a \\
 & \textit{balance} \ (\textit{Node} \ lt \ rt) \\
 & \quad | \ \textit{size} \ rt \ \geq \ 4 * \ \textit{size} \ lt = \textit{balanceToL} \ lt \ rt \\
 & \quad | \ \textit{size} \ lt \ \geq \ 4 * \ \textit{size} \ rt = \textit{balanceToR} \ lt \ rt \\
 & \quad | \ \textit{otherwise} \qquad \qquad \qquad = \textit{Node} \ lt \ rt
 \end{aligned}$$

When testing this function randomly, for example using a *standard generator* for a Quickcheck property [9], the case when the tree is already balanced according to the above definition is tested far more often than the interesting case

when the tree needs balancing. By replacing the branch of the guard requiring the tree to be right-heavy with a target expression, indicated by \bullet , we create a Reach problem which will generate input trees that require balancing.

$$\begin{aligned}
& \textit{balance} :: \textit{Tree } a \rightarrow \textit{Tree } a \\
& \textit{balance } (\textit{Leaf } a) = \textit{Leaf } a \\
& \textit{balance } (\textit{Node } lt \textit{ rt}) \\
& \quad | \textit{size } rt \geq 4 * \textit{size } lt = \bullet \\
& \quad | \textit{size } lt \geq 4 * \textit{size } rt = \textit{balanceToR } lt \textit{ rt} \\
& \quad | \textit{otherwise} = \textit{Node } lt \textit{ rt}
\end{aligned}$$

A solution to the Reach problem with *balance* as the input function is a tree which satisfies the first guard, such as the following:

$$\textit{Node } (\textit{Leaf } 0) (\textit{Node } (\textit{Node } (\textit{Leaf } 1) (\textit{Leaf } 2)) (\textit{Node } (\textit{Leaf } 5) (\textit{Leaf } 2)))$$

This tree can then be used as an input to the original *balance* function to ensure that the auxiliary function *balanceToL* is executed as part of testing. In a similar manner, we can move the target expression to the second branch of the guard to find a tree which ensures that *balanceToR* is executed.

2.1 Forward Reach

In this section we introduce the primary reach solver, Forward Reach, defined by Naylor and Runciman [12, 13]. Forward Reach uses *lazy narrowing* in order to generate inputs efficiently. Lazy narrowing is a concept from functional logic programming [2, 7] and can be described as the natural extension of a non-strict semantics to a language with free variables. Free variables are only bound when their value is required for evaluation to proceed.

To illustrate, we give an example of a lazy narrowing reach solver in action. We show the first steps of an analysis of the *balance* function from the previous section. Each state during evaluation is given by an expression and a *substitution*, a mapping which is an accumulation of the free variable bindings up to the current point of evaluation. For our example, the initial expression is *balance x* and the initial substitution is the trivial mapping $x \mapsto x$ from x to itself.

$$\begin{aligned}
& 1) \{ x \mapsto x \} \\
& \quad \textit{balance } x
\end{aligned}$$

Starting with the trivial mapping rather than the traditional empty mapping helps with the formalisation, as discussed further in section 4.1. The first step of evaluation is to inline the definition for *balance x*:

$$\begin{aligned}
& 2) \{ x \mapsto x \} \\
& \quad \mathbf{case } x \mathbf{ of} \\
& \quad \quad \textit{Leaf } a \rightarrow \textit{Leaf } a \\
& \quad \quad \textit{Node } lt \textit{ rt} \rightarrow \dots
\end{aligned}$$

In order for evaluation to continue the value of the free variable x is now required, which necessitates a *narrowing* step. To begin with, the variable is bound to the leaf constructor for trees by refining the substitution to $x \mapsto \text{Leaf } x'$, and updating the expression being evaluated accordingly.

```

3) {  $x \mapsto \text{Leaf } x'$  }
   case  $\text{Leaf } x'$  of
      $\text{Leaf } a \rightarrow \text{Leaf } a$ 
      $\text{Node } lt \ rt \rightarrow \dots$ 

```

Note that the introduction of a new variable x' is not strictly necessary above, as the manner in which substitutions are later formalised ensures that the variables on the left and right sides of a substitution are independent. Hence, we could equally well use the substitution $x \mapsto \text{Leaf } x$ above, and indeed this simpler approach – which avoids the need to generate a fresh variable name – is used in our subsequent formalisation in section 5. Now that the form of the expression is known, we can reduce the case expression itself:

```

4) {  $x \mapsto \text{Leaf } x'$  }
    $\text{Leaf } x'$ 

```

Evaluation of this execution path terminates with the value $\text{Leaf } x'$. In this case, the target has not been evaluated so the input $\text{Leaf } x'$ is not a solution to the reach problem, independent of any value substituted for x' . Evaluation now backtracks and x is bound to the node constructor for trees. After the narrowing step and following reduction of the case expression we have:

```

5) {  $x \mapsto \text{Node } x_l \ x_r$  }
   if  $\text{size } x_l \geq 4 * \text{size } x_r$  then • else ...

```

Analysis will continue with evaluation of the expression $\text{size } x_l \geq 4 * \text{size } x_r$. Inputs that evaluate to the target will be collected and evaluation will continue until a set number of solutions is found or a given termination condition is reached, e.g. the input has been enumerated to a particular depth.

Lazy narrowing has two key efficiency benefits over the naive approach in which possible inputs are enumerated and evaluated from the beginning each time. First of all, and most importantly, it allows for portions of the input domain to be discarded or accepted if the evaluation concludes while there are still free variables in the substitution, as the same conclusion can be drawn for any input formed by replacing these free variables. This can greatly reduce the search space. For example, above we were able to discard any input of the form $\text{Leaf } x'$. Secondly, some evaluation is shared between different inputs if they have common structure. In particular, their evaluation is shared up to the point where their differences cause execution to take separate branches.

3 A Minimal Language

In this section we introduce the minimal language that we will use for the rest of paper. The language is not suitable for actual programming, but does provide

enough structure to describe the key mechanisms of lazy narrowing. To this end the language has only one type, Peano natural numbers, which provides the simplest example type for showing the recursive mechanics of narrowing. The grammar for expressions of the language is defined as follows:

$$\begin{array}{l}
 \text{Exp} ::= \text{Zero} \\
 \quad | \text{Suc } \text{Exp} \\
 \quad | \bullet \\
 \quad | \text{case } \text{Exp} \text{ of } \text{Exp } \text{Alt} \\
 \quad | \text{var } \text{Var} \\
 \text{Alt} ::= \text{Suc } \text{Var} \rightarrow \text{Exp} \\
 \text{Val} ::= \text{Zero} \mid \text{Suc } \text{Val}
 \end{array}$$

That is, an expression is either a natural number, a target expression \bullet , a case expression, or a variable from some given set Var of names. Case expressions have the form **case** e **of** e_0 f , where the first alternative is the **Zero** branch and the second alternative is the **Suc** branch, which can depend on its argument variable. Expressions are assumed to be closed; variables only appear within the case expression in which they are bound. The values of the language are simply the natural numbers. We do not regard the target expression itself as a value, because our intended interpretation is that the values are ‘normal’ results.

Note that the language does not contain functions or recursion, as these are not required to study the ‘essence’ of lazy narrowing. We do however provide an additional Agda formalisation that incorporates these features, as discussed in section 8. One might also ask why the target expression, which is specific to the Reach problem, is already included in the above language. The reason is simply for convenience: if the target expression was excluded we would need to extend both the syntax and semantics when we later define the Reach problem, whereas including it here means that we only need to extend the syntax.

The behaviour of expressions is defined as a small-step operational semantics, $\rightarrow \subseteq \text{Exp} \times \text{Exp}$, by means of the following inference rules:

$$\begin{array}{c}
 \frac{}{\text{case } \bullet \text{ of } e_0 \text{ } f \rightarrow \bullet} \text{TARGET} \quad \frac{}{\text{case } \text{Zero} \text{ of } e_0 \text{ } f \rightarrow e_0} \text{CASE-Z} \\
 \\
 \frac{}{\text{case } \text{Suc } e \text{ of } e_0 \text{ } (\text{Suc } v \rightarrow e') \rightarrow e'[v := e]} \text{CASE-SUC} \\
 \\
 \frac{e \rightarrow e'}{\text{case } e \text{ of } e_0 \text{ } f \rightarrow \text{case } e' \text{ of } e_0 \text{ } f} \text{SUBJ}
 \end{array}$$

Using a small-step semantics enforces a clear order of evaluation, and supports a natural extension to lazy narrowing. If the case subject is a **Zero** or **Suc** then the semantics are standard, where $e'[v := e]$ denotes the substitution of variable v by the expression e in the expression e' in a capture avoiding manner. The target expression behaves in the same way as an error value, i.e. it is always

propagates through a case expression to the top level, on the basis that once we have found a target no further evaluation is required.

When applying the semantics in practice, we often use the reflexive transitive closure, \rightarrow^* , which is defined in the normal manner:

$$\frac{e \rightarrow e' \quad e' \rightarrow^* e''}{e \rightarrow^* e''} \text{SEQ} \qquad \frac{}{e \rightarrow^* e} \text{REFL}$$

The semantics can be shown by standard methods to be normalising (always terminates in a finite number of steps) and deterministic (always produces a single possible result). However, neither property is a requirement for the definition of the Reach problem or the correctness result which follows.

4 Adding Free Variables

To specify the Reach problem we require a notion of *free* variables. One possibility is to simply allow our expressions to be open, letting the existing variables be free. Although this is the approach taken in the original Reach work [12, 13], we choose to syntactically separate the free variables as an extension of the language. Our reason for making this choice is that free variables are independent of the normal variables of a language; for example, it is easy to make a similar extension to a language that does not have any form of variables.

The extended grammar for expressions is defined below, in which each rule is now parameterised by a set X of free variables, and expressions and values are extended with free variables of the form **fvar** X . Note that we do not require the set of variables for an expression to be minimal, i.e. the set may contain variables that are not used in the expression.

$$\begin{aligned} \text{Exp}_X &::= \text{Zero} \\ &| \text{Suc } \text{Exp}_X \\ &| \bullet \\ &| \text{case } \text{Exp}_X \text{ of } \text{Exp}_X \text{ Alt}_X \\ &| \text{var } \text{Var} \\ &| \text{fvar } X \\ \text{Alt}_X &::= \text{Suc } \text{Var} \rightarrow \text{Exp}_X \\ \text{Val}_X &::= \text{Zero} \mid \text{Suc } \text{Val}_X \mid \text{fvar } X \end{aligned}$$

We will view values of type Val_X as *partial values*, in the sense that they may contain undefined components represented by the free variables. We can also view the original grammars as special cases of the free variable versions in which the free variable sets are empty, i.e. $\text{Exp} \equiv \text{Exp}_\emptyset$, $\text{Alt} \equiv \text{Alt}_\emptyset$ and $\text{Val} \equiv \text{Val}_\emptyset$.

4.1 Substitutions

An input to an expression is a mapping from its free variables to values. In order to define this formally, we first make a slight detour to introduce the more

general notation of a substitution, which will be used later in lazy narrowing. A *substitution* of type $X \rightarrow Y$ is a mapping from the set of free variables X to partial values that contain free variables from the set Y :

$$Sub_{X \rightarrow Y} = X \rightarrow Val_Y$$

Defining substitutions in this manner rather than as a partial mapping from an infinite set of variables results in a simpler formalisation in Agda. In particular, incorporating the set of variables for the domain and range directly into the type removes the need to add the variable sets as constraints later on. A second benefit of this approach is that it yields a monadic interpretation to the composition of substitutions. Given this representation the traditional empty map becomes the trivial map in which each variable is mapped to itself.

Using our notion of substitution, an *input* to an expression can then be viewed as a special case when the set of free variables in the result is empty:

$$Inp_X = Sub_{X \rightarrow \emptyset}$$

We denote substitutions by σ and inputs by τ . The process of applying a substitution is defined recursively in the normal way:

$$\begin{array}{ll} - [-] & :: Exp_X \rightarrow Sub_{X \rightarrow Y} \rightarrow Exp_Y \\ \mathbf{Zero} [\sigma] & = \mathbf{Zero} \\ \mathbf{Suc} e [\sigma] & = \mathbf{Suc} (e [\sigma]) \\ \bullet [\sigma] & = \bullet \\ \mathbf{case} e \mathbf{of} e_0 (\mathbf{Suc} v \rightarrow e') [\sigma] & = \mathbf{case} e [\sigma] \mathbf{of} e_0 [\sigma] (\mathbf{Suc} v \rightarrow e' [\sigma]) \\ \mathbf{var} v [\sigma] & = \mathbf{var} v \\ \mathbf{fvar} x [\sigma] & = \sigma x \end{array}$$

4.2 Reachability

We can now specify the meaning of *reachability* within our framework. Given an expression $e \in Exp_X$ with free variables X , the set of inputs $reach(e) \subseteq Inp_X$ that reach the target expression is defined as follows:

$$\tau \in reach(e) \iff e[\tau] \rightarrow^* \bullet$$

That is, an input τ that provides values for the free variables in expression e satisfies the reachability condition *iff* the input applied to the expression evaluates to the target. This equivalence describes what it means for a given input to reach the target, but does not describe a specific reach problem. An example for such problem might be to find a specific input that satisfies reachability, or to show that none exists. In most languages, but not in our minimal language, the problem is undecidable and therefore an additional termination criterion is included, e.g. find a solution up to a given search depth.

A naive approach to implementing a reach solver is to search for a solution by brute force enumeration and evaluation of all possible inputs. Clearly, however, this is not very efficient. Instead, Naylor and Runciman [12] implement an approach based on lazy narrowing which proves far more efficient. This approach shares evaluation, where possible, across the input domain.

5 Lazy Narrowing Semantics

In this section we define a semantics for our minimal language extended with free variables, based upon the notion of *lazy narrowing*, a symbolic evaluation strategy from functional logic programming. As illustrated in section 2, the basic idea of lazy narrowing is that when evaluation of an expression is suspended on the value of a free variable, we allow evaluation to proceed by performing a *narrowing step*, in which each partial value that the variable could have is considered in turn. As evaluation proceeds a substitution is gradually built up which tracks the instantiation of free variables.

5.1 Preliminaries

We begin by defining a number of concepts that are used in our formalisation of the notion of lazy narrowing, in the form of suspended expressions, minimal narrowing sets, and the composition of substitutions.

Suspended expressions An expression e is *suspended* on a free variable x , denoted by $e \multimap x$, if the value of the variable is required for evaluation of the expression to proceed any further. For our language, the relation $\multimap \subseteq Exp_X \times X$ can be defined by the following two inference rules:

$$\frac{}{\mathbf{fvar} \ x \ \multimap \ x} \text{SUSP} \qquad \frac{e \ \multimap \ x}{\mathbf{case} \ e \ \mathbf{of} \ e_0 \ f \ \multimap \ x} \text{SUBJ-SUSP}$$

That is, free variables are themselves suspended, and a case expression is suspended if its subject expression is suspended. Expressions that are suspended can make no further transitions in our small-step operational semantics from section 3. However, the converse is not true. In particular, values and the target expression cannot make further transitions, but are not suspended.

Minimal narrowing set When an expression is suspended there is a set of possible narrowing steps that can be performed. However, in order to maximise laziness, each of the steps that are considered should be *minimal*, in the sense that it should only instantiate the free variable just enough to allow evaluation to continue, and no further. For our language, in which the only values are natural numbers, this means replacing a free variable x by either **Zero** or **Suc** ($\mathbf{fvar} \ x$), the two possible forms that a natural number can have.

To formalise this idea, we begin by writing x/a for the one-point substitution that maps the free variable $x \in X$ to the partial value $a \in Val_Y$ and leaves all other variables in X unchanged, defined as follows:

$$\begin{array}{l} (/) \\ (x/a) \ x' \mid x \equiv x' \\ \quad \mid \textit{otherwise} = \mathbf{fvar} \ x' \end{array} \quad :: (x \in X) \rightarrow Val_Y \rightarrow Sub_{X \rightarrow X[x/Y]}$$

The return type of the substitution is given by $X [x / Y] = (X - \{x\}) \cup Y$, in which the element $x \in X$ is replaced by the set Y . Note that the type of $(/)$ depends on the name of the variable x , i.e. the operator has a *dependent* type. Being precise in this manner helps to simplify our Agda formalisation. Using this operator we can now define the *minimal narrowing set* $Narr_X(x)$ of a free variable $x \in X$ by replacing x by the two possible forms that it may have:

$$Narr_X(x) = \{x/\text{Zero}, x/\text{Suc}(\text{fvar } x)\}$$

This set has two properties that play an important role in *completeness* of the lazy narrowing semantics. Firstly, the minimal narrowing set itself obeys a notion of completeness, in the sense that for every input that is possible before the narrowing there exists a substitution in which the input remains possible. And secondly, each substitution in the minimal narrowing set is *advancing*, in that it always instantiates a variable. These properties are formalised in section 6.2.

Composition of Substitutions As evaluation proceeds under lazy narrowing, we will construct a substitution in a compositional manner from smaller components. In order to define a composition operator for substitutions, we first note that *Val* forms a monad under the following definitions:

$$\begin{aligned} \text{return} & \quad :: X \rightarrow \text{Val}_X \\ \text{return} & \quad = \text{fvar} \\ (\gg) & \quad :: \text{Val}_X \rightarrow (X \rightarrow \text{Val}_Y) \rightarrow \text{Val}_Y \\ \text{Zero} \gg \sigma & \quad = \text{Zero} \\ \text{Suc } e \gg \sigma & \quad = \text{Suc}(e \gg \sigma) \\ \text{fvar } x \gg \sigma & \quad = \sigma x \end{aligned}$$

We note in passing that this is the *free monad* of the underlying functor for the natural numbers. Using the \gg operator for this monad it is then straightforward to define the composition operator for substitutions:

$$\begin{aligned} (\gg\gg) & \quad :: \text{Sub}_{X \rightarrow Y} \rightarrow \text{Sub}_{Y \rightarrow Z} \rightarrow \text{Sub}_{X \rightarrow Z} \\ \sigma \gg\gg \sigma' & \quad = \lambda a \rightarrow \sigma a \gg \sigma' \end{aligned}$$

Moreover, expanding out the definition of *Sub* in the type for the $\gg\gg$ operator gives $(X \rightarrow \text{Val } Y) \rightarrow (Y \rightarrow \text{Val } Z) \rightarrow (X \rightarrow \text{Val } Z)$, which corresponds to the standard notion of *Kleisli composition* for the *Val* monad.

Along with the monad laws we require one more law, relating the composition of substitutions to the application of a substitution.

Lemma 1. *The sequential application of substitutions to an expression is equivalent to the application of the composed substitutions to the expression:*

$$e[\sigma][\sigma'] \equiv e[\sigma \gg\gg \sigma']$$

5.2 Semantics

We now have all the ingredients required to define a lazy narrowing semantics for our minimal language. A step in the new semantics is either:

- a single step in the original semantics; or
- a minimal narrowing step, if the expression is suspended.

To keep track of the substitutions that are applied during narrowing, we write $e \rightsquigarrow \langle e', \sigma \rangle$ to mean that expression e can make the transition to expression e' in a single step, where σ is the substitution that has been applied in the case of a narrowing step. In the case of a step in the original semantics, we simply return the identity substitution, which is given by the *return* operator of the *Val* monad. More formally, we define a transition relation $\rightsquigarrow \subseteq \text{Exp}_X \times (\text{Exp}_Y \times \text{Sub}_{X \rightarrow Y})$ for lazy narrowing by the following two inference rules:

$$\frac{e \rightarrow_X e'}{e \rightsquigarrow \langle e', \text{return} \rangle} \text{PROM} \qquad \frac{e \dashv\circ x \quad \sigma \in \text{Narr}_X(x)}{e \rightsquigarrow \langle e[\sigma], \sigma \rangle} \text{NARR}$$

The first rule promotes transitions from the original semantics to the new semantics, where $\rightarrow_X \subseteq \text{Exp}_X \times \text{Exp}_X$ is the trivial lifting of the transition relation $\rightarrow \subseteq \text{Exp} \times \text{Exp}$ to operate on expressions with free variables in the set X , for which the inference rules remain syntactically the same as previously except that they now operate on expressions of a more general form. The second rule applies a minimal narrowing step to a suspended expression.

The definition of how to sequence steps in our extended semantics, which takes into account the additional presence of substitutions, is given by a relation \rightsquigarrow^+ that is defined by the following two rules:

$$\frac{e \rightsquigarrow \langle e', \sigma \rangle \quad e' \rightsquigarrow^+ \langle e'', \tau \rangle}{e \rightsquigarrow^+ \langle e'', \sigma \gg \tau \rangle} \text{SEQ} \qquad \frac{e \in \text{Exp}_X \quad \tau \in \text{Inp}_X}{e \rightsquigarrow^+ \langle e[\tau], \tau \rangle} \text{FILL}$$

The first rule simply composes the substitutions from the two component reductions. The second rule adds a final narrowing step to the end of a reduction sequence that instantiates any remaining free variables. The reason for including a final narrowing step is that it simplifies both the definition of forward reachability and its relationship to the original semantics.

5.3 Forward Reachability

Finally, we can now give an alternative characterisation of reachability using our lazy narrowing semantics. Given an expression $e \in \text{Exp}_X$, the set of inputs $\text{reach}_F(e) \in \text{Inp}_X$ that reach the target expression is defined as follows:

$$\tau \in \text{reach}_F(e) \iff e \rightsquigarrow^+ \langle \bullet, \tau \rangle$$

That is, an input τ satisfies the forward reachability condition *iff* there is a lazy narrowing reduction sequence that ends with the target and the given input. The

key difference with our original definition of reachability in section 4.2 is that our new semantics *constructs* an input substitution during the reduction sequence, whereas the original semantics requires that we are *given* a substitution so that it can be applied prior to starting the reduction process. In the next section we show that these two notions of reachability coincide.

6 Correctness of the Narrowing Semantics

To prove that forward reachability is equivalent to the original definition, we first formalise the relationship between the lazy narrowing semantics and the original semantics. This relationship is characterised by two properties, *soundness* and *completeness*, which are proved using a number of lemmas. The proofs of the lemmas themselves are provided in the associated Agda formalisation.

6.1 Soundness

Lemma 2. *A transition in the original semantics can be lifted through a substitution. Given a substitution $\sigma \in \text{Sub}_{X \rightarrow Y}$, we have:*

$$e \rightarrow_X e' \implies e[\sigma] \rightarrow_Y e'[\sigma]$$

Theorem 1 (Soundness). *For every reduction sequence in the lazy narrowing semantics there is a corresponding sequence in the original semantics:*

$$e \rightsquigarrow^+ \langle e', \tau \rangle \implies e[\tau] \rightarrow^* e'$$

Proof. The proof proceeds by rule induction on the definition for the narrowing relation \rightsquigarrow^+ , for which there are three cases to consider.

Case 1 In the base case when the narrowing is a simple application of

$$\frac{}{e \rightsquigarrow^+ \langle e[\tau], \tau \rangle} \text{ FILL}$$

the goal follows immediately from the reflexivity of \rightarrow^* :

$$\frac{}{e[\tau] \rightarrow^* e[\tau]} \text{ REFL}$$

Case 2 There are two inductive cases to consider, depending on the nature of the first reduction in a narrowing sequence. We first consider the case when the reduction is a narrowing step, constructed as follows:

$$\text{NARR} \frac{\frac{e \multimap x \quad \sigma \in \text{Narr}_X(x)}{e \rightsquigarrow \langle e[\sigma], \sigma \rangle} \quad e[\sigma] \rightsquigarrow^+ \langle e', \tau \rangle}{e \rightsquigarrow^+ \langle e', \sigma \rightsquigarrow \tau \rangle} \text{ SEQ}$$

We are now free to use the three assumptions $e \multimap x$, $\sigma \in \text{Narr}_X(x)$ and $e[\sigma] \rightsquigarrow^+ \langle e', \tau \rangle$ in our proof. In this case, we only require the third of these assumptions in order to verify our goal, by first using the induction hypothesis (IH) $e[\sigma] \rightsquigarrow^+ \langle e', \tau \rangle \implies e[\sigma][\tau] \rightarrow^* e'$, and then applying lemma 1:

$$\frac{\frac{e[\sigma] \rightsquigarrow^+ \langle e', \tau \rangle}{e[\sigma][\tau] \rightarrow^* e'} \text{ IH}}{e[\sigma] \gg \tau \rightarrow^* e'} \text{ LEMMA 1}$$

Case 3 We now consider the case when the first reduction is a promoted reduction from the original language, constructed as follows:

$$\text{PROM} \frac{e \rightarrow_X e'}{e \rightsquigarrow \langle e', \text{return} \rangle} \quad \frac{e' \rightsquigarrow^+ \langle e'', \tau \rangle}{e \rightsquigarrow^+ \langle e'', \text{return} \gg \tau \rangle} \text{ SEQ}$$

In this case our goal can then be verified by lifting the reduction from the original language through the input substitution using lemma 2, sequencing with the result of applying the induction hypothesis to the remaining reduction sequence, and finally applying an identity law for Kleisli composition:

$$\text{LEMMA 2} \frac{\frac{e \rightarrow_X e'}{e[\tau] \rightarrow e'[\tau]} \quad \frac{e' \rightsquigarrow^+ \langle e'', \tau \rangle}{e'[\tau] \rightarrow^* e''} \text{ IH}}{e[\tau] \rightarrow^* e''} \text{ SEQ}}{e[\text{return} \gg \tau] \rightarrow^* e''} \text{ ID}$$

□

Although the above proof was presented specifically for the specific case of lazy narrowing semantics, it is not dependent on the properties of the narrowing set or the condition for applying a narrowing step. Therefore the proof is also valid for any narrowing set and any applicability condition.

6.2 Completeness

Definition 1. We exploit two pre-orderings on substitutions, which respectively capture the idea of one substitution being a *prefix* or *suffix* of another:

$$\sigma_1 \sqsubseteq \sigma_2 \iff \exists \sigma'. \sigma_1 \gg \sigma' \equiv \sigma_2$$

$$\sigma_1 \leq \sigma_2 \iff \exists \sigma'. \sigma' \gg \sigma_1 \equiv \sigma_2$$

Lemma 3. *If the source expression of a transition in the original semantics is not suspended then the transition can be ‘unlifted’. Given a substitution $\sigma \in \text{Sub}_{X \rightarrow Y}$ and a transition $e[\sigma] \rightarrow_Y e'$ for which $e \not\multimap x$, we have:*

$$\exists e'_\sigma. e \rightarrow_X e'_\sigma \wedge e'_\sigma[\sigma] \equiv e'$$

Lemma 4. *The lazy narrowing set is complete. For every input there is a substitution in the narrowing set that is a prefix of the input:*

$$\forall x \in X, \tau \in \text{Inp}_X. \exists \sigma \in \text{Narr}_X(x). \sigma \sqsubseteq \tau$$

Lemma 5. *The lazy narrowing set is advancing. The identity substitution is a strict prefix of every substitution in the narrowing set:*

$$\forall x \in X, \sigma \in \text{Narr}_X(x). \text{return} \sqsubset \sigma$$

Theorem 2 (Completeness). *For every reduction sequence in the original semantics there is a corresponding reduction in the lazy narrowing semantics:*

$$e[\tau] \rightarrow^* e' \implies e \rightsquigarrow^+ \langle e', \tau \rangle$$

Proof. The proof proceeds by rule induction on the definition for the evaluation relation \rightarrow^* , for which there are three cases to consider.

Case 1 In the base case when the evaluation is just reflexivity

$$\frac{}{e[\tau] \rightarrow^* e[\tau]} \text{REFL}$$

the goal follows immediately by instantiating free variables:

$$\frac{}{e \rightsquigarrow^+ \langle e[\tau], \tau \rangle} \text{FILL}$$

Case 2 There are two inductive cases to consider, depending on whether or not the expression e is suspended when the sequencing rule is applied:

$$\frac{e[\tau] \rightarrow e' \quad e' \rightarrow^* e''}{e[\tau] \rightarrow^* e''} \text{SEQ}$$

In the case when e is not suspended our goal can be verified as follows, in which the two branches of the proof tree exploit the two conclusions from lemma 3:

$$\frac{\text{PROM} \frac{\text{LEMMA 3} \frac{}{e \rightarrow e'_\tau}}{e \rightsquigarrow \langle e'_\tau, \text{return} \rangle} \quad \frac{\text{LEMMA 3} \frac{e' \rightarrow^* e''}{e'_\tau[\tau] \rightarrow^* e''}}{e'_\tau \rightsquigarrow^+ \langle e'', \tau \rangle} \text{IH}}{\frac{}{e \rightsquigarrow^+ \langle e'', \text{return} \gg \tau \rangle} \text{SEQ}}{\frac{}{e \rightsquigarrow^+ \langle e'', \tau \rangle} \text{ID}}$$

Case 3 Finally, when e is suspended on x , because the narrowing set $\text{Narr}(x)$ is complete (lemma 4) there is a substitution in this set that is a prefix of the input τ , i.e. a substitution $\sigma \in \text{Narr}(x)$ and input τ' for which $\tau \equiv \sigma \gg \tau'$. Based upon this observation our goal can then be verified as follows:

$$\begin{array}{c}
\text{NARR} \frac{e \multimap x \quad \sigma \in \text{Narr}(x)}{e \rightsquigarrow \langle e[\sigma], \sigma \rangle} \quad \frac{\frac{e[\tau] \rightarrow^* e'}{e[\sigma][\tau'] \rightarrow^* e'} \text{ LEMMA 1}}{e[\sigma] \rightsquigarrow^+ \langle e', \tau' \rangle} \text{ IH}}{e \rightsquigarrow^+ \langle e', \sigma \gg \tau' \rangle} \text{ SEQ} \\
\hline
e \rightsquigarrow^+ \langle e', \tau \rangle \quad \text{LEMMA 4}
\end{array}$$

Well-foundedness In the third case above, we need to explicitly verify that the induction is well-founded as the induction hypothesis is not trivially smaller in this case. Instead, with each iteration the *input* gets smaller. To formalise this well-foundedness neatly and generally, we restrict our notion of substitutions $\text{Sub}_{X \rightarrow Y}$ to the case when the free variable sets X and Y are finite, and every variable in Y appears in the result of the substitution. For our purposes this leads to no loss of generality and all of our definitions satisfy these restrictions. With these in place, we then have the following two results, which together with lemma 5 ensures that the use of induction in the third case is well-founded.

Lemma 6. *The suffix relation $<$ is well-founded. For any substitution τ_0 , there only exists finite chains of substitutions τ_i such that:*

$$\tau_n < \dots < \tau_1 < \tau_0$$

Lemma 7. *A suffix formed by an advancing prefix is strict.*

$$\sigma \gg \sigma_1 \equiv \sigma_2 \wedge \text{return} \sqsubset \sigma \implies \sigma_1 < \sigma_2$$

□

Whereas the soundness proof was independent of the properties of the narrowing set and the condition for its applicability, the completeness proof relies on the fact that the narrowing set is complete and advancing, and that narrowing steps can always be applied when an expression is suspended.

6.3 Correctness

Using the soundness and completeness results, it is now straightforward to prove that our two notions of reachability are equivalent:

Theorem 3 (Correctness). *For all expressions $e \in \text{Exp}_X$:*

$$\text{reach}_F(e) \equiv \text{reach}(e)$$

Proof.

$$\begin{aligned}
\tau \in \text{reach}_F(e) &\iff e \rightsquigarrow^+ \langle \bullet, \tau \rangle && \text{(by definition)} \\
&\iff e[\tau] \rightarrow^* \bullet && \text{(theorems 1 and 2)} \\
&\iff \tau \in \text{reach}(e) && \text{(by definition)}
\end{aligned}$$

7 Agda Formalisation

Our correctness result has also been formalised in the Agda [14]. The Agda formalisation follows the presentation given in the paper closely: the language grammar and semantic rules convert directly to inductive datatypes, and rule induction translates to recursive dependent functions. A proof of the main result and all associated lemmas is available online from:

<http://tinyurl.com/reachtheory>

Using Agda brings a number of important benefits. First of all, it provides a guarantee that our results are correct. Secondly, it helped guide the development of our theory and proofs, resulting in a number of simplifications. For example, when translating our original formalisation into Agda we found that it contained a subtle error. The process of correcting the error also pointed towards a neater theory. In particular, our original lazy narrowing formulation kept the substitution as an environment, only replacing free variables when they were needed. The most natural way to fix the error was to apply the substitution to the current expression immediately, removing the need to keep the substitution as an environment. This also removed an unnecessary distinction in the formalisation: in the original formulation the expression/environment pair $\langle e, \sigma \rangle$ behaved equivalently to the pair $\langle e[\sigma], \sigma \rangle$, yet the two were distinct. And finally, the use of Agda had a positive effect on the formulation of the representation of substitutions. In order to ensure totality in Agda we had to parameterise substitutions with the set of variables used in their domain and result. Far from being a hindrance, this led us to the monadic formulation of composition.

8 Extending the language

In this paper we focused on a minimal language to emphasise the key elements of the reach problem and a solver based on lazy narrowing semantics. However, our results also scale up to a more realistic language that includes function application, lambda abstraction and fixed points [4]. This section briefly describes the changes that are required to the Agda formalisation.

First of all, the expression grammar is extended to include the three new constructors: function application, lambda abstraction and fixed points. To avoid ill-formed expressions the addition of these language features requires the new expression grammar to be typed. Therefore a function type is added to the language, along with the type of natural numbers. The small step semantics is extended to account for the new language constructs.

Our formalisation of the lazy narrowing semantics for the extended language restricts free variables, and by extension narrowing, to natural numbers. Although this is certainly a limitation, it is standard in the lazy narrowing literature, where a narrowing theory is generally described for first-order data initially, and then potentially extended to the higher-order case in subsequent work. With this restriction, the alteration to the lazy narrowing semantics and

correctness proof is minor. The suspension predicate, $e \multimap x$, has to be updated as an expression can now be suspended within a function application or a fix-point expression. We defined the lazy narrowing semantics by lifting the original semantics, and this definition remains unchanged except that we now lift the extended semantics. Finally, the lemmas, particularly the lift and unlift lemmas (2&3), need updating to account for the additional cases. The proof of soundness and completeness remain identical under the updated lemmas.

The ease of this extension suggests it may be possible to generalise the theory by abstracting away from the details of the underlying language and semantics that is used, which is an interesting topic for further work.

9 Related Work

There is a large body of work on the theory of lazy narrowing in functional logic programming. We introduce and compare two particularly relevant theories to ours. In their seminal work, Antoy et al. [2] established the soundness and completeness of the related notion of needed narrowing, and the optimality of needed narrowing within a restricted domain. However, whereas our formalisation is based on extending a small-step semantics, theirs is based on classical rewrite systems. As a result, our approach is easier to mechanically verify, which we have done, as the semantics of our language has a direct representation in proof assistants. In fact, to the best of our knowledge, this is a first time that a lazy narrowing formalisation has had such a verification.

A formulation of lazy narrowing which is more closely related to ours is given by Albert et al. [1] in which a “natural” big-step semantics is defined before an implementation driven small-step semantics is introduced. Both semantics are call-by-need, implement sharing, and are proved to be equivalent. They go on to extend the small-step semantics with additional features such as equational constraints and external functions. There is a difference in motive in comparison to our work, as they establish lazy narrowing as a programming language feature whereas we are interested in using lazy narrowing to analyse the operation of a program. The difference manifests itself in the theories: they relate their small-step semantics back to their defining big-step semantics, whereas we relate our lazy narrowing semantics back to the underlying functional semantics.

10 Conclusions and Future Work

In this article we established the correctness of a reach solver for a minimal language, based upon a soundness and completeness result for a lazy narrowing semantics. Our final formulation of the semantics is the result of several iterations and improvements, and captures the main ideas of lazy narrowing in a simple and concise manner. In particular, the use of an underlying small-step semantics was instrumental in simplifying the theory. The simplicity along with the use of precise types enables a direct translation of our result to the Agda system [4].

There are number of interesting directions in which the theory developed in this article could be extended and improved, which are summarised below.

Other reach solvers The work in this article lays the ground for attempting to formalise alternative and more general reach solvers, such as the Backward Reach solver defined in Naylor’s thesis [13]. In addition, tools such as Lindblad’s data generator [10] and Lazy SmallCheck [16] define logical *or* operators that evaluate both argument expressions in parallel, which could significantly improve the performance of lazy narrowing as expressions of the form $e \text{ or } e'$ can be reduced to true if either argument reduces to true in the current substitution state. We could easily add such an operator to our language. However, our formulation suggests a generalisation to this idea, in the form of evaluating branches in parallel and utilising equational reasoning on case expressions.

Other language features We used a minimal language for simplicity, but it is important to consider how our approach generalises to other language features. For algebraic datatypes, we expect it should be straightforward to extend our theory using ideas from generic programming as in [8], while first-order functions could be handled by representing functions using tries as in the improved Lazy Smallcheck [15]. Another interesting area to explore is dependent type theory. Lazy narrowing is often used in automated property based testing and dependent type theory seems a natural coupling as it offers an inbuilt language for specifications. In this area there is also potential for interesting comparison to related work such as automated proof search [14].

Efficiency We showed that the lazy narrowing definition of reachability for our language is correct with respect to the original specification of reachability. However we have not made any formal argument regarding the *efficiency* of the lazy narrowing approach, either against an alternative narrowing semantics or a naive approach based on brute force search. Such an argument could be made on the basis of simply counting the number of reduction steps required, or adopt a more sophisticated approach, for example using the idea of *improvement theory* [11], which has recently been used to prove that a general purpose optimisation technique for lazy languages never makes programs worse [6].

Acknowledgements

We would like to thank members of the FP Lab in Nottingham and the anonymous referees for useful comments and suggestions regarding this work.

References

- [1] Albert, E., Hanus, M., Frank, H., Oliver, J., Germán, V.: Operational semantics for declarative multi-paradigm languages. *Journal of Symbolic Computation* 40(1), 795–829 (2005)

- [2] Antoy, S., Echahed, R., Hanus, M.: A Needed Narrowing Strategy. *Journal of the ACM* 47(4), pp. 776–822 (2000)
- [3] Claessen, K., Hughes, J.: QuickCheck: A Lightweight Tool for Random Testing of Haskell Programs. In: *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming* (2000)
- [4] Fowler, J.: *Towards a Theory of Reach - Agda Proof* (2015), <https://github.com/JonFowler/theoryofreach>
- [5] Gill, A., Runciman, C.: Haskell Program Coverage. In: *Proceedings of the ACM SIGPLAN Workshop on Haskell* (2007)
- [6] Hackett, J., Hutton, G.: Worker/Wrapper/Makes It/Faster. In: *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming* (2014)
- [7] Hanus, M.: A Unified Computation Model for Functional and Logic Programming. In: *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (1997)
- [8] Hinze, R.: A New Approach to Generic Functional Programming. In: *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (2000)
- [9] Hughes, J.: QuickCheck: An Automatic Testing Tool for Haskell (QuickCheck manual), <http://www.cse.chalmers.se/~rjmh/QuickCheck/manual.html>
- [10] Lindblad, F.: Property Directed Generation of First-Order Test Data. In: *Proceedings of the Eighth Symposium on the Trends in Functional Programming* (2007)
- [11] Moran, A., Sands, D.: Improvement in a Lazy Context: An Operational Theory for Call-by-need. In: *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (1999)
- [12] Naylor, M., Runciman, C.: Finding Inputs that Reach a Target Expression. In: *Proceedings of the 7th IEEE International Working Conference on Source Code Analysis and Manipulation* (2007)
- [13] Naylor, M.F.: *Hardware-Assisted and Target-Directed Evaluation of Functional Programs*. Ph.D. thesis, University of York (2008)
- [14] Norell, U.: *Towards a Practical Programming Language Based on Dependent Type Theory*. Ph.D. thesis, Goteborg University (2007)
- [15] Reich, J.S., Naylor, M., Runciman, C.: Advances in Lazy SmallCheck. In: *Proceedings of the 24th Symposium on the Implementation and Application of Functional Languages* (2013)
- [16] Runciman, C., Naylor, M., Lindblad, F.: SmallCheck and Lazy SmallCheck Automatic Exhaustive Testing for Small Values. In: *Proceedings of the First ACM SIGPLAN Symposium on Haskell* (2008)