

A Semantics for Context-Sensitive Reduction Semantics

Casey Klein¹, Jay McCarthy², Steven Jaconette¹, and Robert Bruce Findler¹

¹ Northwestern University

² Brigham Young University

Abstract. This paper explores the semantics of the meta-notation used in the style of operational semantics introduced by Felleisen and Hieb. Specifically, it defines a formal system that gives precise meanings to the notions of contexts, decomposition, and plugging (recomposition) left implicit in most expositions. This semantics is not naturally algorithmic, so the paper also provides an algorithm and proves a correspondence with the declarative definition.

The motivation for this investigation is PLT Redex, a domain-specific programming language designed to support Felleisen-Hieb-style semantics. This style of semantics is the de-facto standard in operational semantics and, as such, is widely used. Accordingly, our goal is that Redex programs should, as much as possible, look and behave like those semantics. Since Redex’s first public release more than seven years ago, its precise interpretation of contexts has changed several times, as we repeatedly encountered reduction systems that did not behave according to their authors’ intent. This paper describes the culmination of that experience. To the best of our knowledge, the semantics given here accommodates even the most complex uses of contexts available.

1 Introduction

The dominant style of operational semantics in use today has at its heart the notion of a context that controls where evaluation occurs. These contexts allow the designer of a reduction semantics to factor the definition of a calculus into one part that specifies the atomic steps of computation and a second part that controls where these steps may occur. This factoring enables concise specification, e.g., that a language is call-by-value or call-by-name or call-by-need (Ariola and Felleisen 1997), that `if` expressions must evaluate the test position before the branches, and even that exceptions, continuations, and state (Felleisen and Hieb 1992) behave in the expected ways, all without cluttering the rules that describe the atomic steps of computation.

Unfortunately, the precise meaning of context decomposition has not been nailed down in a way that captures its diverse usage in the literature. Although an intuitive definition is easy to understand from a few examples, this intuition does not cover the full power of contexts. For example, which terms match the pattern $C[e]$ from this language, in which values and contexts are mutually referential?

$$\begin{aligned} C &::= [] \mid (v C) \mid (C e) \\ e &::= v \mid x \mid (e e) \\ v &::= (\lambda (x) e) \mid (\text{cont } C) \end{aligned}$$

And which terms match this bizarre, small language?

$$C ::= C[(f \ _)] \mid _$$

To remedy this lack, we have developed a semantics for matching and reduction that not only supports these exotic languages but also captures the intuitive meanings of countless existing research papers. This semantics does not assume explicit language-specific definitions of plugging and decomposition, since most expositions leave these concepts implicit.

Our motivation for studying context-sensitive matching is its implementation in the domain-specific programming language Redex (Felleisen et al. 2010; Matthews et al. 2004). Redex is designed to support the semantics engineer with a lightweight toolset for operational semantics and related formalisms. Specifically, Redex supports rapid prototyping of context-sensitive operational semantics, random testing, automatic type-setting, and, via its embedding in Racket, access to a large palette of standard programming tools. Redex is widely used, having supported several dozen research papers as well as the latest Scheme standard (Sperber et al. 2007) and a number of larger models, including one of the Racket virtual machine (Klein et al. 2010).

In keeping with the spirit of Redex, we augment a standard proof-based validation of our work with testing. More concretely, in addition to proving a correspondence between a specification of context-sensitive matching and an algorithm for that specification, we have conducted extensive testing of the semantics, using a Redex model of Redex (there is little danger of meta-circularity causing problems, as the embedding uses a modest subset of Redex’s functionality—notably, no contexts or reduction relations). This model allows us to test that our semantics gives the intended meanings to interesting calculi from the literature, something that would be difficult to prove.

The remainder of this paper builds up an intuitive understanding of what contexts are and how they are used via a series of examples, gives a semantics for Redex’s rewriting system, and discusses an algorithm to implement the semantics.

2 Matching and Contexts

This section introduces the notion of contexts and explains through a series of examples how matching works in their presence. Each example comes with a lesson that informs the design of our context-sensitive reduction semantics semantics.

In its essence, a pattern of the form $C[e]$ matches an expression when the expression can be split into two parts, an outer part (the context) that matches C and an inner part that matches e . The outer part marks where the inner part appears with a hole, written $_$. In other words, when thinking of an expression as a tree, matching against $C[e]$ finds some subtree of the expression that matches e , and then replaces that sub-term with a hole to build a new expression in such a way that the new expression matches C .

$$\begin{aligned} a &::= (+ \ a \ a) \mid \text{number} \\ C &::= (+ \ C \ a) \mid (+ \ a \ C) \mid _ \\ C[(+ \ \text{number}_1 \ \text{number}_2)] &\longrightarrow \\ C[\Sigma[_[\text{number}_1, \text{number}_2]]] & \end{aligned}$$

Figure 1: Arithmetic Expressions

To get warmed up, consider figure 1. In this language a matches addition expressions and C matches contexts for addition expressions. More precisely, C matches an addition expression that has exactly one hole. For example, the expression $(+ 1 2)$ matches $C[a]$ three ways, as shown in figure 2. Accordingly, the reduction relation given in figure 1 reduces addition expressions wherever they appear in an expression, e.g., reducing $(+ (+ 1 2) (+ 3 4))$ to two different expressions, $(+ 3 (+ 3 4))$ and $(+ (+ 1 2) 7)$. This example tells us that our context matching semantics must support multiple decompositions for any given term.

A common use of contexts is to restrict the places where reduction may occur in order to model a realistic programming language's order of evaluation. Figure 3 gives a definition of E that enforces call-by-value left-to-right order of evaluation. For example, consider this nested set of function calls, $((f x) (g y))$, in which the

result of $(g y)$ is passed to the result of $(f x)$. It decomposes into the context $([] (g y))$, allowing evaluation in the first position of the application. It does not, however, decompose into the context $((f x) [])$, since the grammar for E allows the hole to appear in the argument position of an application expression only when the function position is already a value. Accordingly, the reduction system insists that the call to f happens before the call to g . This example tells us that our semantics for decomposition must be able to support multiple different ways to decompose each expression form, depending on the subexpressions of that form (application expressions in this case).

Contexts can also be used in clever ways to model the call-by-need λ -calculus. Like call-by-name, call-by-need evaluates the argument to a function only if the value is actually needed by the function's body. Unlike call-by-name, each function argument is evaluated at most once. A typical implementation of a language with call-by-need uses state to track if an argument has been evaluated, but it is also possible to give a direct explanation, exploiting contexts to control where evaluation occurs.

Figure 4 shows the contexts from Ariola and Felleisen (1997)'s model of call-by-need. The first three of E 's alternatives are standard, allowing evaluation in the argument of the $+1$ primitive, as well as in the function position of an application (regardless of what appears in the argument position). The fourth alternative allows

evaluation in the body of a λ -expression that is in the function position of an application. Intuitively, this case says that once we have determined the function to be applied, then we can begin to evaluate its body. Of course, the function may eventually need its argument, and at that point, the final alternative comes into play. It says that when an applied function needs its argument, then that argument may be evaluated.

$$\begin{array}{ll} C = [] & a = (+ 1 2) \\ C = (+ [] 2) & a = 1 \\ C = (+ 1 []) & a = 2 \end{array}$$

Figure 2: Example Decomposition

$$\begin{array}{l} e ::= (e e) \mid x \mid v \\ v ::= (\lambda (x) e) \mid +1 \mid number \\ E ::= (E e) \mid (v E) \mid [] \end{array}$$

Figure 3: λ -calculus

$$\begin{array}{l} E ::= [] \mid (+1 E) \mid (E e) \\ \quad \mid ((\lambda (x) E) e) \\ \quad \mid ((\lambda (x) E[x]) E) \end{array}$$

Figure 4: Call-by-need Contexts

As an example, the expression $((\lambda (x) (+1\ 1)) (+1\ 2))$ reduces by simplifying the body of the λ -expression to 2 , without reducing the argument, because it decomposes into this context $((\lambda (x) []) (+1\ 2))$ using the fourth alternative of E . In contrast, $((\lambda (x) (+1\ x)) (+1\ 2))$ reduces to $((\lambda (x) (+1\ x))\ 3)$ because the body of the λ -expression decomposes into the context $(+1\ [])$ with x in the hole, and thus the entire expression decomposes into the context $((\lambda (x) (+1\ x)) [])$. This use of contexts tells us that our semantics must be able to support a sophisticated form of nesting, namely that sometimes a decomposition must occur in one part of a term in order for a decomposition to occur in another.

When building a model of first-class continuations, there is an easy connection to make, namely that an evaluation context is itself a natural representation for a continuation. That is, at the point that a continuation is grabbed, the context in which it is grabbed is the continuation. Figure 5 extends the left-to-right call-by-value model in figure 3 with support for continuations. It adds **call/cc**, the operator that grabs a continuation, and the new value form **(cont E)** that represents a continuation and can be applied to invoke the continuation.

$$\begin{array}{l}
 v ::= \dots \mid \text{call/cc} \mid (\text{cont } E) \\
 E[(\text{call/cc } v)] \longrightarrow E[(v (\text{cont } E))] \\
 E_1[(\text{cont } E_2) v] \longrightarrow E_2[v]
 \end{array}$$

Figure 5: Continuations

For example, the expression $(+1 (\text{call/cc } (\lambda (k) (k\ 2))))$ reduces by grabbing a continuation. In this model that continuation is represented as $(\text{cont } (+1\ []))$, which is then applied to **call/cc**'s argument in the original context, yielding the expression $(+1 ((\lambda (k) (k\ 2)) (\text{cont } (+1\ []))))$. The next step is to substitute for **k**, which yields the expression $(+1 ((\text{cont } (+1\ []))\ 2))$. This expression has a continuation value in the function position of an application, making the next step invoke the continuation. So, we can simply replace the context of the continuation invocation with the context inside the continuation, plugging the argument passed to the continuation in the hole, yielding $(+1\ 2)$. This reduction system tells us that our context decomposition semantics must be able to support contexts that appear in a term that play no part in any decomposition (and yet must still match a specified pattern, such as E).

Generalizing from ordinary continuations to delimited continuations is simply a matter of factoring the contexts into two parts, one that may contain prompts and one that may not. Figure 6 shows one way to do this, as an extension of the call-by-value lambda calculus from figure 3.

$$\begin{array}{l}
 e ::= \dots \mid (\# e) \\
 v ::= \dots \mid \text{call/comp} \mid (\text{comp } M) \\
 M ::= (M e) \mid (v M) \mid [] \\
 E ::= M \mid E[(\# M)] \\
 E[(\# M[(\text{call/comp } v)])] \longrightarrow \\
 E[(\# M[(v (\text{comp } M))])]
 \end{array}$$

Figure 6: Delimited Continuations

The non-terminal E matches an arbitrary evaluation context and M matches an evaluation context that does not contain any prompt expressions. Accordingly, the rule for grabbing a continuation exploits this factoring to record only the portion of the context between the call to **call/comp** and the nearest enclosing prompt.

the portion of the context between the call to **call/comp** and the nearest enclosing prompt.

The interesting aspect of this system is how E refers to M and how that makes it difficult to support an algorithm that matches E . For all of the example systems in this section so far, a matching algorithm can match a pattern of the form $C[e]$ by attempting to match C against the entire term and, once a match has been found, attempting to match what appeared at the hole against e . With E , however, this leads to an infinite loop because E expands to a decomposition that includes E in the first position.³

A simple fix that works for the delimited continuations example is to backtrack when encountering such cycles; that fix, however, does not work for the first definition of C given in figure 7. Specifically, C would match only \square with an algorithm that treats that cycle as a failure to match, but the context $(f \square)$ should match C , and more generally, the two definitions of C in figure 7 should be equivalent.

$$C ::= C[(f \square)] \mid \square$$

$$C ::= (f C) \mid \square$$

Figure 7: Wacky Context

3 A Semantics for Matching

This section formalizes the notion of matching used in the definitions of the example reduction systems in section 2. For ease of presentation, we stick to the core language of patterns and terms in figure 8. Redex supports a richer language of patterns (notably including a notion of Kleene star), but this core captures an essence suitable for explaining the semantics of matching.

Ignoring embedded contexts, a term t is simply a binary tree where leaf nodes are atoms a and interior nodes are constructed with **cons**. A context C is similarly a binary tree, but with a distinguished path (marked with **left** and **right**) from the root of the context to its **hole**.

Contexts are generated by decomposition and represent single-holed contexts. Although the **hole** can appear multiple times in a single term, such terms represent expressions that contain multiple, independently pluggable contexts.

Patterns p take one of six forms. Atomic patterns a and the **hole** pattern match only themselves. A pattern (**name** $x p$) binds the pattern variable x to the term matched by p . Repeated pattern variables force the corresponding sub-terms to be identical. A pattern (**nt** n) matches terms that match any of the alternatives of the non-terminal n (defined outside the pattern). We write decomposition patterns $p_1[p_2]$ using a separate keyword for clarity: (**in-hole** $p_1 p_2$). Finally, interior nodes are matched by the pattern (**cons** $p_1 p_2$), where p_1 and p_2 match the corresponding sub-terms.

$$t ::= (\text{cons } t t)$$

$$\quad \mid a$$

$$\quad \mid C$$

$$C ::= \text{hole}$$

$$\quad \mid (\text{left } C t)$$

$$\quad \mid (\text{right } t C)$$

$$p ::= a$$

$$\quad \mid (\text{name } x p)$$

$$\quad \mid (\text{nt } n)$$

$$\quad \mid (\text{in-hole } p p)$$

$$\quad \mid (\text{cons } p p)$$

$$\quad \mid \text{hole}$$

$$a \in \text{Literals}$$

$$x \in \text{Variables}$$

$$n \in \text{Non-Terminals}$$

Figure 8: Patterns and Terms

³ Some find the equivalent, non-problematic grammar $E ::= M \mid M[(\# E)]$ clearer. At least one author of the present paper (who has spent a considerable amount of time hacking on Redex's implementation, no less), however, does not and was surprised when Redex failed to terminate on a similar example. We have also received comments from Redex users who were surprised by similar examples, suggesting that Redex should support such definitions.

For example, the left-hand side of the reduction rule in figure 1 corresponds to the following pattern, where the literal **empty** is used for the empty sequence and the pattern **number** matches literal numbers:

$$\begin{aligned} &(\text{in-hole } (\text{name } C \text{ (nt } C)) \\ &\quad (\text{cons } + \\ &\quad\quad (\text{cons } (\text{name } \textit{number}_1 \text{ number}) \\ &\quad\quad\quad (\text{cons } (\text{name } \textit{number}_2 \text{ number}) \\ &\quad\quad\quad\quad \text{empty})))) \end{aligned}$$

Figure 9 gives a semantics for patterns via the judgment form $G \vdash t : p \mid b$, which defines when the pattern p matches the term t . The grammar G is a finite map from non-terminals to sets of patterns. The bindings b is a finite map from pattern variables to terms showing how the pattern variables of p can be instantiated to yield t . The $G \vdash t : p \mid b$ judgment relies on an auxiliary judgment $G \vdash t = C[t'] : p \mid b$ that performs decompositions. Specifically, it holds when t can be decomposed into a context C that matches p and contains the sub-term t' at its hole.

Many of the rules for these two judgment forms rely on the operator \sqcup . It combines two mappings into a single one by taking the union of their bindings, as long as the domains do not overlap. If the domains do overlap, then the corresponding ranges must be the same; otherwise \sqcup is not defined. Accordingly, rules that use \sqcup apply only when \sqcup is well-defined.

The $G \vdash t : p \mid b$ rules are organized by the structure of p . The atom and hole rules produce an empty binding map because those pattern contain no pattern variables. The **name** rule matches p with t and produces a map extended with the binding (x, t) . The **nt** rule applies if any of the non-terminal's alternatives match. The scope of an alternative's pattern variables is limited to that alternative, and consequently, the **nt** rule produces an empty binding map. The **cons** rule matches the sub-terms and combines the resulting sets of bindings. The **in-hole** rule uses the decomposition judgment form to find a decomposition and checks that the term in the hole matches p_2 .

The rules for the $G \vdash t = C[t'] : p \mid b$ form are also organized around the pattern. The **hole** decomposition rule decomposes any term t into the empty context and t itself. The first of two **cons** decomposition rules applies when a decomposition's focus may be placed within a pair's left sub-term. This decomposition highlights the same sub-term t'_i as the decomposition of t_i does, but places it within the larger context (**left** $C t_2$). The second of the **cons** decomposition rules does the same for the pair's right sub-term. The **nt** decomposition rule propagates decompositions but, as in the corresponding matching rule, ignores binding maps.

The **in-hole** decomposition rule performs a nested decomposition. Nested decomposition occurs, for example, when decomposing according to call-by-need evaluation contexts (see the last production in figure 4). The **in-hole** rule decomposes t into a composed context $C_1 ++ C_2$ and a sub-term t' , where p_1 and p_2 match C_1 and C_2 respectively. The definition of context composition (figure 9, bottom-right) follows the path in C_1 . The **name** decomposition rule is similar to the corresponding matching rule, but it introduces a binding to the context that is matched, not the entire term.

$$\boxed{G \vdash t : p \mid b} \quad \overline{G \vdash a : a \mid \emptyset} \quad \overline{G \vdash \text{hole} : \text{hole} \mid \emptyset}$$

$$\frac{G \vdash t : p \mid b}{G \vdash t : (\text{name } x \ p) \mid \{(x, t)\} \sqcup b}$$

$$\frac{p \in G(n) \quad G \vdash t : p \mid b}{G \vdash t : (\text{nt } n) \mid \emptyset}$$

$$\frac{G \vdash t_1 : p_1 \mid b_1 \quad G \vdash t_2 : p_2 \mid b_2 \quad k \in \{\text{cons, left, right}\}}{G \vdash (k \ t_1 \ t_2) : (\text{cons } p_1 \ p_2) \mid b_1 \sqcup b_2}$$

$$\frac{G \vdash t_1 = C[t_2] : p_1 \mid b_1 \quad G \vdash t_2 : p_2 \mid b_2}{G \vdash t_1 : (\text{in-hole } p_1 \ p_2) \mid b_1 \sqcup b_2}$$

$$\boxed{G \vdash t = C[t'] : p \mid b}$$

$$\overline{G \vdash t = \text{hole}[t] : \text{hole} \mid \emptyset}$$

$$\frac{G \vdash t_1 = C[t'_1] : p_1 \mid b_1 \quad G \vdash t_2 : p_2 \mid b_2 \quad k \in \{\text{cons, left, right}\}}{G \vdash (k \ t_1 \ t_2) = (\text{left } C \ t_2)[t'_1] : (\text{cons } p_1 \ p_2) \mid b_1 \sqcup b_2}$$

$$\frac{G \vdash t_1 : p_1 \mid b_1 \quad G \vdash t_2 = C[t'_2] : p_2 \mid b_2 \quad k \in \{\text{cons, left, right}\}}{G \vdash (k \ t_1 \ t_2) = (\text{right } t_1 \ C)[t'_2] : (\text{cons } p_1 \ p_2) \mid b_1 \sqcup b_2}$$

$$\frac{p \in G(n) \quad G \vdash t_1 = C[t_2] : p \mid b}{G \vdash t_1 = C[t_2] : (\text{nt } n) \mid \emptyset}$$

$$\frac{G \vdash t = C_1[t_1] : p_1 \mid b_1 \quad G \vdash t_1 = C_2[t_2] : p_2 \mid b_2}{G \vdash t = (C_1 \ ++ \ C_2)[t_2] : (\text{in-hole } p_1 \ p_2) \mid b_1 \sqcup b_2}$$

$$\frac{G \vdash t_1 = C[t_2] : p \mid b}{G \vdash t_1 = C[t_2] : (\text{name } x \ p) \mid \{(x, C)\} \sqcup b}$$

$$\begin{array}{l}
G \in \text{Non-Terminal} \rightarrow \wp(p) \\
b \in \text{Variable} \rightarrow t
\end{array}
\quad
\begin{array}{l}
\text{hole } ++ \ C \quad = \ C \\
(\text{left } C_1 \ t) ++ \ C_2 \quad = \ (\text{left } (C_1 \ ++ \ C_2) \ t) \\
(\text{right } t \ C_1) ++ \ C_2 \quad = \ (\text{right } t \ (C_1 \ ++ \ C_2))
\end{array}$$

Figure 9: Matching and Decomposition

4 An Algorithm for Matching

The rules in figure 9 provide a declarative definition of context-sensitive matching, but they do not lead directly to a tractable matching algorithm. There are two problems. First, as reflected in the two **CONS** decomposition rules, an algorithm cannot know a priori whether to match on the left and decompose on the right or to decompose on the left and match on the right. An implementation that tries both possibilities scales exponentially in the number of nested **CONS** patterns matched (counting indirect nesting through non-terminals). Second, the rules provide no answer to the question of whether to proceed in expanding a non-terminal if none of the input term has been consumed since last encountering that non-terminal. This question arises, for example, when decomposing by the non-terminal **E** from the grammar in figure 6, since **E**'s second alternative causes the **in-hole** rule to decompose the same term by **E**. This second problem is the manifestation of left recursion in the form of grammars we consider.

The first problem can be solved by matching and decomposing simultaneously. Since these tasks differ only in their treatment of **hole** patterns, much work can be saved by sharing intermediate results between the tasks. Figure 10 demonstrates this approach with a function **M** that returns a set of pairs (d, b) representing possible ways to match or decompose the input term. In a pair representing a match, d is the marker *****; in a pair representing a decomposition, d is a pair (C, t) such that the input term can be decomposed into a context C and a sub-term t occurring in C 's hole.

The first two **M** cases handle the pattern **hole**. If the term in question is also **hole**, then it may be considered either to match **hole** or to decompose into **hole** in the empty context. If the term is not **hole**, then only decomposition is possible. The third case handles atomic patterns by producing a match result only if the given term is identical to the atom.

The (meta) context in which a call to **M** appears may eventually discard some or all of the results it receives. For example, consider the fourth clause, which handles **CONS** patterns. If the term is also a pair (constructed with any of **cons**, **left**, or **right**), then this case makes two recursive calls and examines the cross product of the results using the **select** helper function. For each result pair, the case merges their bindings and checks that the results are not both decompositions. If neither is a decomposition, **select** combines the pair into a match result; if exactly one is a decomposition, it extends the decomposition with the term matched by the non-decomposition. If both are decompositions, then the match fails.

The next case, for patterns (**in-hole** $p_c p_h$), recurs with p_c and the input term, expecting to receive decompositions. For each one, it makes another recursive call, this time with p_h and the sub-term in the decomposition's focus. Each of the latter call's results m is combined with the decomposition's context, yielding a match result if m is a match and a larger context if m is a decomposition.

The remaining three cases are straightforward. The **name** case recurs on the sub-pattern and extends the bindings of each of the results with either the matched term or the context carved out by the decomposition. The **nt** case tries each alternative, discarding the binding component of each result. The final case, a catch-all, applies when the pattern does not match or decompose the input term.

$$\begin{aligned} \text{matches} : G \ p \ t \rightarrow \wp(b) & & m ::= (d, b) \\ \text{matches}[[G, p, t]] = \{b \mid (\bullet, b) \in M[[G, p, t]]\} & & d ::= (C, t) \mid \bullet \end{aligned}$$

$$\begin{aligned} M : G \ p \ t \rightarrow \wp(m) \\ M[[G, \text{hole}, \text{hole}]] &= \{((\text{hole}, \text{hole}), \emptyset), (\bullet, \emptyset)\} \\ M[[G, \text{hole}, t]] &= \{((\text{hole}, t), \emptyset)\} \\ M[[G, a, a]] &= \{(\bullet, \emptyset)\} \\ M[[G, (\text{cons } p_l \ p_r), (k \ t_l \ t_r)]] &= \{(d, b) \mid k \in \{\text{cons}, \text{left}, \text{right}\}, \\ &\quad d \in \text{select}[[t_l, d_l, t_r, d_r]], \\ &\quad b = b_l \sqcup b_r, \\ &\quad (d_r, b_r) \in M[[G, p_r, t_r]], \\ &\quad (d_l, b_l) \in M[[G, p_l, t_l]]\} \\ M[[G, (\text{in-hole } p_c \ p_h), t]] &= \{(d, b) \mid d = \text{combine}[[C, d_h]], \\ &\quad b = b_c \sqcup b_h, \\ &\quad (d_h, b_h) \in M[[G, p_h, t_c]], \\ &\quad ((C, t_c), b_c) \in M[[G, p_c, t]]\} \\ M[[G, (\text{name } x \ p), t]] &= \{(d, b) \mid b' = \{(x, \text{named}[[d, t]])\} \sqcup b, \\ &\quad (d, b) \in M[[G, p, t]]\} \\ M[[G, (\text{nt } n), t]] &= \{(d, \emptyset) \mid (d, b) \in M[[G, p, t]], p \in G(n)\} \\ M[[G, p, t]] &= \{\} \end{aligned}$$

$$\begin{aligned} \text{select} : t \ d \ t \ d \rightarrow \wp(d) \\ \text{select}[[t_l, \bullet, t_2, \bullet]] &= \{\bullet\} \\ \text{select}[[t, (C, t_1), t_2, \bullet]] &= \{((\text{left } C \ t_2), t_1)\} \\ \text{select}[[t_l, \bullet, t_2, (C, t_2)]] &= \{((\text{right } t_l \ C), t_2)\} \\ \text{select}[[t_l, (C, t_1), t_2, (C', t_2)]] &= \{\} \end{aligned}$$

$$\begin{aligned} \text{combine} : C \ d \rightarrow d \\ \text{combine}[[C, \bullet]] &= \bullet \\ \text{combine}[[C_1, (C_2, t)]] &= (C_1 \ ++ \ C_2, t) \end{aligned}$$

$$\begin{aligned} \text{named} : d \ t \rightarrow t \\ \text{named}[[\bullet, t]] &= t \\ \text{named}[[C, (t_1), t_2]] &= C \end{aligned}$$

Figure 10: Core matching algorithm (cases apply in order)

Putting aside the problem of left recursion, the call $M[[G, p, t]]$ computes the set of b such that $G \vdash t : p \mid b$ or $G \vdash t = C[t'] : p \mid b$ for some C and t' , and the top-level wrapper function **matches** restricts this set to the bindings associated with match derivations.

To make this precise, we first give a definition of left-recursion. Intuitively, a grammar is left-recursive if there is a way, in a straight-forward recursive parser, to get from some non-terminal back to that same non-terminal without consuming any input. So, our definition of left-recursion builds a graph from the grammar by connecting each pattern to the other patterns that might be reached without consuming any input, and then checks for a cycle in the graph. The most interesting case is the last one, where an **in-hole** pattern is connected to its second argument when the first argument can generate **hole**.

Definition. A grammar G is *left recursive* if $p \rightarrow_G^* p$ for some p , where \rightarrow_G^* is the transitive (but not reflexive) closure of $\rightarrow_G \subseteq p \times p$, the least relation satisfying the following conditions:

1. (**nt** n) $\rightarrow_G p$ if $p \in G(n)$,
2. (**name** $n p$) $\rightarrow_G p$.
3. (**in-hole** $p p'$) $\rightarrow_G p$,
4. (**in-hole** $p p'$) $\rightarrow_G p'$ if $G \vdash \mathbf{hole} : p \mid b$, and

Theorem. For all G, p , and t , if G is not left recursive, then $b \in \mathbf{matches}[[G, p, t]] \Leftrightarrow G \vdash t : p \mid b$.

The complete proof is available at eecs.northwestern.edu/~robby/plugin/.

Parsing algorithms that support left recursive context-free grammars go back nearly fifty years (Kuno 1965). We refer the reader to Frost et al. (2007, section 3) for a summary. We have implemented an extension of the packrat parsing algorithm (Warth et al. 2008) that dynamically detects left recursion and treats the choice leading to it as a failure. If the other choices for the same portion of the input make any progress at all, the algorithm repeats the parse attempt, in hopes that the entries added to the memo table during the failed attempt will cause a second attempt to succeed. This process continues as long as repeated attempts make additional progress. Extending the algorithm in figure 10 with a similar iterative phase allows matching of terms from left recursive grammars, such the ones in figure 6 and figure 7.

5 A Semantics for Reduction

We now put the notion of matching from section 3 to work in a formalization of the standard notation for context-sensitive reduction rules. As with patterns, we consider a core specification language that lacks many of the conveniences of a language like Redex but nevertheless highlights the principal ideas.

Figure 11 shows our definition. A user of Redex specifies a grammar and rules of the shape $p \longrightarrow r$, each consisting of a pattern p and a term template r . Redex uses the judgment form in the upper-left corner of the figure to determine if a particular term t

$$\frac{G \vdash t : p \mid b \quad t' = \text{inst}[[r, b]]}{G \vdash t / p \longrightarrow t' / r}$$

$$r ::= a$$

- | hole
- | (var x)
- | (app $f r$)
- | (in-hole $r r$)
- | (cons $r r$)

$$f \in t \rightarrow t$$

$$\begin{aligned} \text{inst} : r b \rightarrow t \\ \text{inst}[[a, b]] &= a \\ \text{inst}[[\text{hole}, b]] &= \text{hole} \\ \text{inst}[[\text{var } x, b]] &= b(x) \\ \text{inst}[[\text{in-hole } r_1 r_2, b]] &= \text{plug}[[\text{inst}[[r_1, b]], \text{inst}[[r_2, b]]]] \\ \text{inst}[[\text{cons } r_1 r_2, b]] &= \text{join}[[\text{inst}[[r_1, b]], \text{inst}[[r_2, b]]]] \\ \text{inst}[[\text{app } f r, b]] &= \delta(f, \text{inst}[[r, b]]) \end{aligned}$$

$$\begin{aligned} \text{plug} : C t \rightarrow t \\ \text{plug}[[\text{hole}, t]] &= t \\ \text{plug}[[\text{left } C_l t_r, C]] &= (\text{left } \text{plug}[[C_l, C]] t_r) \\ \text{plug}[[\text{left } C_l t_r, t]] &= (\text{cons } \text{plug}[[C_l, t]] t_r) \\ \text{plug}[[\text{right } t_l C_r, C]] &= (\text{right } t_l \text{plug}[[C_r, C]]) \\ \text{plug}[[\text{right } t_l C_r, t]] &= (\text{cons } t_l \text{plug}[[C_r, t]]) \end{aligned}$$

$$\begin{aligned} \text{join} : t t \rightarrow t \\ \text{join}[[C, t]] &= (\text{left } C t) \quad \text{where no-ctxts } t \\ \text{join}[[t, C]] &= (\text{right } t C) \quad \text{where no-ctxts } t \\ \text{join}[[t_1, t_2]] &= (\text{cons } t_1 t_2) \end{aligned}$$

$\delta : (t \rightarrow t) t \rightarrow t$
An unspecified function that applies metafunctions

$$\frac{}{\text{no-ctxts } a} \quad \frac{\text{no-ctxts } t_1 \quad \text{no-ctxts } t_2}{\text{no-ctxts } (\text{cons } t_1 t_2)}$$

Figure 11: A semantics for reduction (function cases apply in order)

reduces to t' by the given rule. The grammar in the figure's top-right gives the syntax for term templates, which include atoms, the context **hole**, references to variables bound by the left-hand side, applications of meta-level functions (e.g., substitution), hole-filling operations, and pairing operations.

The rest of the figure defines template instantiation. Atoms and **hole** instantiate to themselves, variables instantiate to their values, and meta-applications instantiate to the result of applying the meta-function to the instantiated argument template.

The instantiation of **in-hole** templates makes use of a generic **plug** function that accepts a context and a term and returns the result of plugging the context with the term. When **plug**'s second argument is a context, it constructs a larger context by concatenating the two contexts, preserving the path to the hole. The path extension is necessary, for example, to support the following rule for an unusual control operator:

$$E[(\text{call/cc2 } v)] \longrightarrow E[(v (\text{cont } E[E]))]$$

When **plug**'s second argument is some non-context term, it replaces the **left** or **right** constructor with **CONS**, producing a non-context term.

Insistence that **plug**'s first argument be a context creates a potential problem for rules which extend contexts, like this one for another unusual control operator:

$$E[(\text{call/cc+ } v)] \longrightarrow E[(v (\text{cont } (+1 E)))]$$

Although the rule does not explicitly define a path for the extended context $(+1 E)$, one can be safely inferred, since the term paired with E has no pluggable sub-terms.

The case of the **inst** function for **CONS** templates performs this inference via the function **join**. When given a context and a term containing no contexts, **join** extends the context's path through the extra layer. When both arguments contain contexts, **join** combines the terms with **CONS**, preventing possible ambiguity in a subsequent plugging operation.

Note, however, that the embedded contexts themselves remain pluggable by reduction rules and meta-functions that later pick apart the result term. For example, consider the rule for yet another unusual control operator:

$$E[(\text{call/ccs } v)] \longrightarrow E[(v (\text{tuple } (\text{cont } E) (\text{cont } E[E])))]$$

This rule calls v with a pair of continuation values. The term denoting this pair is not itself pluggable, but the embedded contexts can be plugged by subsequent reduction steps, after they are extracted by the reduction rules for projecting **tuple** components.

In addition to these contrived reduction rules, the semantics in figure 11 supports all of the systems in section 2, as well as the most sophisticated uses of contexts we have encountered in the literature, specifically:

- Ariola and Felleisen (1997)'s core call-by-need calculus. Their extension of this calculus to **letrec** uses decomposition in fundamentally the same way, but the particular formulation they choose makes use of pattern-matching constructs orthogonal to the ones we describe here, namely associative-commutative matching and a Kleene star-like construct that enforces dependencies between adjacent terms.

The examples directory distributed with Redex shows one way to define their `le-trec` evaluation contexts without these constructs, which Redex does not currently support.

- Flatt et al. (2007)’s semantics for delimited control in the presence of dynamic binding, exception handling, and Scheme’s `dynamic-wind` form.
- Chang and Felleisen (2011)’s call-by-need calculus, which defines evaluation contexts using a heavily left-recursive grammar.

6 Related Work

Barendregt (1984) makes frequent use of a notion of contexts specialized to λ -terms. Like ours, these contexts may contain multiple holes, but plug’s behavior differs in that it fills all of the context’s holes. Felleisen and Hieb (1992) exploit the power of a selective notion of context to give equational semantics for many aspects of programming languages, notably continuations and state. The meaning of multi-holed grammars does not arise in their work, since the grammar for contexts restricts them to exactly one hole.

Lucas (1995) later explored an alternative formulation of selective contexts. This formulation defines contexts not by grammars but by specifying, for each function symbol, which sub-term positions may be reduced. Because the specification depends only on the function symbol’s identity (i.e., and not on its sub-terms), this formulation cannot express common evaluation strategies, such as left-to-right, call-by-value evaluation. Follow-up work on this form of context-sensitive rewriting focuses on tools for proving termination, generally a topic of limited interest when studying reduction systems designed to model a programming language since these systems are not expected to terminate.

As part of their work on SL, a meta-language similar to Redex, Xiao et al. (2001) define a semantics for Felleisen-Hieb contexts by translating grammars into their own formalism, an extension of finite tree automata. This indirect approach allows SL to prove decomposition lemmas automatically using existing automata algorithms, but it is considerably more complicated than our approach and does not allow for multi-holed grammars like the ones in figure 5 and figure 6.

Dubois (2000) develops the first formulation of a Felleisen-Hieb reduction semantics in a proof assistant, as part of a mechanized proof of the soundness of ML’s type system. Her formulation encodes single-hole contexts as meta-level term-to-term functions (restricted to coincide with the usual grammar defining call-by-value evaluation) and therefore models plug as meta-application. The formulation does not use an explicit notion of decomposition; instead, the contextual closure reduction rule applies to terms that may be formed using the plug operation.

Berghofer’s, Leroy’s, and Xi’s solutions to the POPLmark Challenge (Aydemir et al. 2005) use Dubois’s encoding for the challenge’s reduction semantics. Vouillon’s solution uses a first-order encoding of contexts and therefore provides an explicit definition of plugging. The other submitted solutions use structural operational semantics, do not address dynamic semantics at all, or are no longer available online.

Danvy and Nielsen (2004) and Sieczkowski et al. (2010) provide an axiomatization of the various components of a Felleisen-Hieb reduction semantics, such as a decomposition relation, that together define the semantics. This axiomatization is not an appropriate basis for Redex for two reasons. First, it requires users to specify plugging and decomposition explicitly. Common practice leaves these definitions implicit, and one of our design goals for Redex is to support conventional definitions. Second, the axiomatization requires decomposition to be a (single-valued) function, ruling out the semantics in figure 1 and, more problematically, reduction semantics for multi-threaded programs and programs in languages like C and Scheme, which do not specify an order of evaluation for application expressions.

More broadly speaking, there are hundreds⁴ of papers that use evaluation context semantics to model programming languages for just as many different purposes. Although we have not implemented anywhere near all of them in Redex, we have sought out interesting and non-standard ones over the years to try them out and to build our intuition about how a semantics should behave.

Acknowledgments Thanks to Stephen Chang for his many interesting examples of contexts that challenged our understanding of context-sensitive matching. Thanks also to Matthias Felleisen and Matthew Flatt for helpful discussions of the work.

A version of this paper can be found online at:

<http://www.eecs.northwestern.edu/~robby/plug/>

That web page contains the final version of the paper as it appears in the proceedings and the Redex models for all of the figures in this paper.

Bibliography

Zena M. Ariola and Matthias Felleisen. The Call-by-Need Lambda-Calculus. *J. Functional Programming* 7(3), pp. 265–301, 1997.

Brian E. Aydemir, Aaron Bohannon, Matthew Fairbairn, J. Nathan Foster, Benjamin C. Pierce, Peter Sewell, Dimitrios Vytiniotis, Geoffrey Washburn, Stephanie Weirich, and Steve Zdancewic. Mechanized Metatheory for the Masses: The POPLmark Challenge. In *Proc. Intl. Conf. Theorem Proving in Higher Order Logics*, Lecture Notes in Computer Science volume 3603, pp. 50–65, 2005.

H. P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. North Holland, 1984.

⁴ There are more than 400 citations to the original Felleisen-Hieb paper; while evaluation-context based semantics are still widely used, the paper is now rarely cited as it has become a standard part of the programming languages landscape.

- Stephen Chang and Matthias Felleisen. The Call-by-need Lambda Calculus. Unpublished Manuscript, 2011.
- Olivier Danvy and Lasse R. Nielsen. Refocusing in Reduction Semantics. Aarhus University, BRICS RS-04-26, 2004.
- Catherine Dubois. Proving ML Type Soundness within Coq. In *Proc. Intl. Conf. Theorem Proving in Higher Order Logics*, Lecture Notes in Computer Science volume 1869, pp. 126–144, 2000.
- Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. *Semantics Engineering with PLT Redex*. MIT Press, 2010.
- Matthias Felleisen and Robert Hieb. The Revised Report on the Syntactic Theories of Sequential Control and State. *Theoretical Computer Science* 103(2), pp. 235–271, 1992.
- Matthew Flatt, Gang Yu, Robert Bruce Findler, and Matthias Felleisen. Adding Delimited and Composable Control to a Production Programming Environment. In *Proc. ACM Intl. Conf. Functional Programming*, pp. 165–176, 2007.
- Richard A. Frost, Rahmatullah Hafiz, and Paul C. Callaghan. Modular and Efficient Top-Down Parsing for Ambiguous Left-Recursive Grammars. In *Proc. International Conference on Parsing Technology*, pp. 109–120, 2007.
- Casey Klein, Matthew Flatt, and Robert Bruce Findler. The Racket Virtual Machine and Randomized Testing. 2010. <http://plt.eecs.northwestern.edu/racket-machine/>
- Susumu Kuno. The Predictive Analyzer and a Path Elimination Technique. *Communications of the ACM* 8(7), pp. 453–462, 1965.
- Salvador Lucas. Fundamentals of Context-Sensitive Rewriting. In *Proc. Seminar on Current Trends in Theory and Practice of Informatics*, Lecture Notes in Computer Science volume 1012, pp. 405–412, 1995.
- Jacob Matthews, Robert Bruce Findler, Matthew Flatt, and Matthias Felleisen. A Visual Environment for Developing Context-Sensitive Term Rewriting Systems. In *Proc. Intl. Conf. Rewriting Techniques and Applications*, Lecture Notes in Computer Science volume 3091, pp. 301–311, 2004.
- Filip Sieczkowski, Malgorzata Biernacka, and Dariusz Biernacki. Automating Derivations of Abstract Machines from Reduction Semantics: A Generic Formalization of Refocusing in Coq. In *Proc. Symp. Implementation and Application of Functional Languages*, To appear in Lecture Notes in Computer Science, 2010.
- Michael Sperber, R. Kent Dybvig, Matthew Flatt, Anton van Straaten, Richard Kelsey, William Clinger, Jonathan Rees, Robert Bruce Findler, and Jacob Matthews. Revised [6] Report on the Algorithmic Language Scheme. Cambridge University Press, 2007.
- Alessandro Warth, James R. Douglass, and Todd Millstein. Packrat Parsers Can Support Left Recursion. In *Proc. ACM SIGPLAN Wksp. Partial Evaluation and Program Manipulation*, pp. 103–110, 2008.
- Yong Xiao, Amr Sabry, and Zena M. Ariola. From Syntactic Theories to Interpreters: Automating the Proof of Unique Decomposition. *Higher-Order and Symbolic Computation* 14(4), pp. 387–409, 2001.