Predicate Dispatching: A Unified Theory of Dispatch

Michael Ernst Craig Kaplan Craig Chambers

Technical report UW-CSE-98-01-02
Department of Computer Science and Engineering
University of Washington
Seattle, WA, USA 98195-2350
{mernst,csk,chambers}@cs.washington.edu
http://www.cs.washington.edu/research/projects/cecil/

12 January 1998

Abstract

Predicate dispatching generalizes previous method dispatch mechanisms by permitting arbitrary predicates to control method applicability and by using logical implication between predicates as the overriding relationship. The method selected to handle a message send can depend not just on the classes of the arguments, as in ordinary object-oriented dispatch, but also on the classes of subcomponents, on an argument's state, and on relationships between objects. This simple mechanism subsumes and extends object-oriented single and multiple dispatch, ML-style pattern matching, predicate classes, and classifiers, which can all be regarded as syntactic sugar for predicate dispatching. This paper introduces predicate dispatching, gives motivating examples adapted from a prototype implementation, and presents its static and dynamic semantics.

1 Introduction

Many programming languages support some mechanism for dividing the body of a (generic) function into a set of cases, with a declarative mechanism for selecting the right case for each dynamic invocation of the generic function. Case selection can be broken down into tests for *applicability* (a case is a candidate for invocation if its guard is satisfied) and *overriding* (which selects one of the applicable cases for invocation).

Object-oriented languages use overloaded methods as the cases. A method is applicable if the run-time class of the receiver argument is the same as or a subclass of the class on which the receiver is specialized. Multiple dispatching [BKK$^+$86, Cha92] enables testing the classes of all of the arguments. One method overrides another if its specializer classes are subclasses of the other's, using either lexicographic (CLOS [Ste90]) or pointwise (Cecil [Cha93a]) ordering.

Predicate classes [Cha93b] automatically classify an object of class A as an instance of virtual subclass B (a subclass of A) whenever B's predicate (an arbitrary expression typically testing the runtime state of an object) is true. This creation of virtual class hierarchies makes method dispatching applicable even in cases where the effective class of an object may change over time. Classifiers [HHM90b] and modes [Tai93] are similar mechanisms for reclassifying an object into one of a number of subclasses based on a case-statement-like test of arbitrary boolean conditions.

Pattern matching (as in ML [MTH90]) bases applicability tests on the run-time datatype constructor tags of the arguments and their subcomponents. As with classifiers and modes, textual ordering determines overriding. Some languages, such as Haskell [HJW$^+$92], allow arbitrary boolean guards to accompany patterns, restricting applicability. Views [Wad87] extend pattern matching to abstract data types by enabling them to offer various concrete datatype-like interfaces.

Predicate dispatching integrates, generalizes, and provides a uniform interface to these similar but previously incomparable mechanisms. A method declaration specifies its applicability via a *predicate expression*,

E	$\in$ *expr*	The set of expressions in the underlying programming language
T	$\in$ *type*	The set of types in the underlying programming language
c	$\in$ *class-id*	The namespace of classes
m, f	$\in$ *method-id*	The namespace of methods and fields
p	$\in$ *pred-id*	The namespace of predicate abstractions
v, w	$\in$ *var-id*	The namespace of variables

Figure 1: Syntactic domains and variables

which is a logical formula over class tests (i.e., tests that an object is of a particular class or one of its subclasses) and arbitrary boolean-valued expressions from the underlying programming language. A method is applicable when its predicate expression evaluates to *true*. Method m_1 overrides method m_2 when m_1's predicate logically implies that of m_2; this relationship is computed at compile time. Static typechecking verifies that, for all possible combinations of arguments to a generic function, there is always a single most-specific applicable method. This ensures that there are no message-not-understood errors (called match-not-exhaustive in ML) or message-ambiguous errors at run-time.

Predicate expressions capture the basic primitive mechanisms underlying a wide range of declarative dispatching mechanisms. Combining these primitives in an orthogonal and general manner enables new sorts of dispatching that are not expressible by previous dispatch mechanisms. Predicate dispatching preserves several desirable properties from its object-oriented heritage, including that methods can be declared in any order and that new methods can be added to existing generic functions without modifying the existing methods or clients; these properties are not shared by pattern-matching-based mechanisms.

Section 2 introduces the syntax, semantics, and use of predicate dispatching through a series of examples. Section 3 defines its dynamic and static semantics formally. Section 4 discusses predicate tautology testing, which is the key mechanism required by the dynamic and static semantics. Section 5 surveys related work, and Section 6 concludes with a discussion of future directions for research.

2 Overview

This section demonstrates some of the capabilities of predicate dispatching by way of a series of examples. We incrementally present a high-level syntax which appears in full in Figure 5; Figure 1 lists supporting syntactic domains. Predicate dispatching is parameterized by the syntax and semantics of the host programming language in which predicate dispatching is embedded.

2.1 Dynamic dispatch

Each method implementation has an attached predicate expression which specifies when the method is applicable. Predicate expressions include class tests and negations, conjunctions, and disjunctions of predicate expressions. An omitted predicate expression indicates that its method handles all type-correct arguments.

method-sig	::=	**signature** m **(** $\langle$ T $\rangle$ **)** **:** T	
method-decl	::=	**method** m **(** $\langle$ *formal-pattern* $\rangle$ **)** [**when** *pred-expr*] *method-body*	
pred-expr	::=	*expr* @ c	succeeds if *expr* evaluates to an instance or subclass of class c
	|	**not** *pred-expr*	negation
	|	*pred-expr* **and** *pred-expr*	conjunction (short-circuited)
	|	*pred-expr* **or** *pred-expr*	disjunction (short-circuited)

Method signature declarations give the type signature shared by a family of method implementations. A message send expression need examine only the corresponding method signature declaration to determine its type-correctness, while a set of overloaded method implementations must completely and unambiguously implement the corresponding signature in order to be type-correct.

Predicate dispatching can simulate both singly- and multiply-dispatched methods by *specializing* formal parameters on a class (via the "`@`*class*" syntax). Specialization limits the applicability of a method to objects that are instances of the given class or one of its subclasses. ML-style pattern matching is modeled by considering each constructor of a datatype to be a class and specializing methods on the constructor classes. More generally, predicate dispatching supports the construction of arbitrary conjunctions, disjunctions, and negations of class tests. The following example uses predicate dispatching to implement the `Zip` function which converts a pair of lists into a list of pairs:[1]

```
type List;
  class Cons subtypes List { head:Any, tail:List };
  class Nil subtypes List;

signature Zip(List, List):List;
method Zip(l1, l2) when l1@Cons and l2@Cons {
    return Cons(Pair(l1.head, l2.head), Zip(l1.tail, l2.tail)); }
method Zip(l1, l2) when l1@Nil or l2@Nil { return Nil; }
```

The first `Zip` method tests the classes of both arguments, and it only applies when both are instances of `Cons` (or some subclass); this is an implicit conjunction of two class tests. The second `Zip` method uses explicit disjunction to test whether either argument is an instance of `Nil` (or some subclass). The type checker can verify statically that the two implementations of `Zip` are mutually exclusive and exhaustive over all possible arguments that match the signature, ensuring that there will be no "message not understood" or "message ambiguous" errors at run-time, without requiring the cases to be put in any particular order.

ML-style pattern matching requires all cases to be written in one place and put in a particular total order, resolving ambiguities in favor of the first successfully matching pattern. In a traditional (singly- or multiply-dispatched) object-oriented language without the ability to order cases, either the base case of `Zip` must be written as the default case for all pairs of `List` objects (unnaturally, and unsafely in the face of future additions of new subclasses of the default type), or *three* separate but identical base methods must be written (one for `Nil`×`Any`, one for `Any`×`Nil`, and a third for `Nil`×`Nil` to resolve the ambiguity between the first two). In our experience with object-oriented languages (using a pointwise, not lexicographic, ordering), these triplicate base methods for binary messages occur frequently.

As a syntactic convenience, class tests can be written in the formal argument list:

formal-pattern ::= [v] [@ c]   like v @ c in *pred-expr*

The first `Zip` method above would then be rewritten as

```
method Zip(l1@Cons, l2@Cons) {
  return Cons(Pair(l1.head, l2.head), Zip(l1.tail, l2.tail)); }
```

2.2 Pattern matching

Predicates can test the run-time classes of components of an argument, just as pattern matching can query substructures, by suffixing the `@`*class* test with a record-like list of field names and corresponding class tests; names can be bound to field contents at the same time. The ability in pattern matching to test for particular constants of built-in types is a simple extension of class tests.

pred-expr ::= ...
    expr @ *specializer*
specializer ::= c [**{** ⟨ *field-pat* ⟩ **}**]
field-pat ::= m [= v] [@ *specializer*]

As with pattern matching, testing the representation of components of an object makes sense when the object and the tested components together implement a single abstraction. We do not advocate using pattern matching to test components of objects in a way that crosses natural abstraction boundaries.

---

[1] `Any` is the top class, subclassed by all other classes, and `Pair` returns an object containing its two arguments.

Our syntax for pattern matching on records is analogous to that for creating a record: `{ x := 7, y := 22 }` creates a two-component record, binding the `x` field to 7 and the `y` field to 22, while `{ x = xval }` pattern-matches against a record containing an `x` field, binding the new variable `xval` to the contents of that field and ignoring any other fields that might be present. The similarity between the record construction and matching syntaxes follows ML. Our presentation syntax also uses braces for record type specifiers (as in the declaration of the `Cons` class, above) and to delimit code blocks (as in the definitions of the `Zip` methods, above).

The following example, adapted from our implementation of an optimizing compiler, shows how a `ConstantFold` method can dispatch for binary operators whose arguments are constants and whose operator is integer addition:

```
type Expr;
  signature ConstantFold(Expr):Expr;
  -- default constant-fold optimization: do nothing
  method ConstantFold(e) { return e; }

  type AtomicExpr subtypes Expr;
    class VarRef subtypes AtomicExpr { ... };
    class IntConst subtypes AtomicExpr { value:int };
    ... -- other atomic expressions here

  type Binop;
    class IntPlus subtypes Binop { ... };
    class IntMul subtypes Binop { ... };
    ... -- other binary operators here

  class BinopExpr subtypes Expr { op:Binop, arg1:AtomicExpr, arg2:AtomicExpr, ... };
    -- override default to constant-fold binops with constant arguments
    method ConstantFold(e@BinopExpr{ op@IntPlus, arg1@IntConst, arg2@IntConst }) {
      return new IntConst { value := arg1 + arg2 }; }
    ... -- many more similarly expressed cases for other operators here

  class UnopExpr subtypes Expr { op:Unop, arg:AtomicExpr, ... };
  ...
```

2.3 Boolean expressions

To increase the expressiveness of predicate dispatching, predicates may include arbitrary boolean expressions from the underlying programming language. Additionally, names may be bound to values, for use later in the predicate expressions and in the method body. Expressions from the underlying programming language that appear in predicate expressions should have no externally observable side effects.

pred-expr ::= ...
| **test** *expr* succeeds if *expr* evaluates to *true*
| **let** v **:=** E bind v to E; always succeeds

The following extension to the `ConstantFold` example illustrates these features.

```
-- Handle case of adding zero to anything (but don't be ambiguous
-- with existing method for zero plus a constant).
method ConstantFold(e@BinopExpr{ op@IntPlus, arg1@IntConst{ value=v }, arg2=a2 })
  when test(v == 0) and not(a2@IntConst) {
  return a2; }
method ConstantFold(e@BinopExpr{ op@IntPlus, arg1=a1, arg2@IntConst{ value=v } })
  when test(v == 0) and not(a1@IntConst) {
```

```
    return a1; }

  ... -- other special cases for operations on 0,1 here
```

2.4 Predicate abstractions

Named predicate abstractions can factor out recurring tests and give names to semantically meaningful concepts in the application domain. To allow abstraction over both tests and variable bindings, predicate abstractions can return a record-like list of bindings. These bindings resemble the fields of a record or class, and similar support is given to pattern matching against a subset of the results of a named predicate invocation. Predicate abstractions thus can act like views or virtual subclasses of some object (or tuple of objects), with the results of predicate abstractions acting like the virtual fields of the virtual class. If the properties of an object tested by a collection of predicates are mutable, the object may be given different virtual subclass bindings at different times in its life, providing the benefits of using classes to organize code even in situations where an object's "class" is not fixed.

Because object identity is not affected by these different views on an object, named predicate abstractions are more flexible than coercions in environments with side-effects. A single object can be classified in multiple independent ways by different predicate abstractions without being forced to define all the possible conjunctions of independent predicates as explicit classes, relieving some of the problems associated with a mix-in style of class organization [HHM90b, HHM90a].

pred-sig ::= **predsignature** *p* **(** ⟨ *T* ⟩ **)** **return** **{** ⟨ *f* **:** *T* ⟩ **}**
pred-decl ::= **predicate** *p* **(** ⟨ *formal-pattern* ⟩ **)**
    [**when** *pred-expr*] [**return** **{** ⟨ *f* **:=** *expr* ⟩ **}**]
pred-expr ::= . . .
    | *p* **(** ⟨ *expr* ⟩ **)** [**=>** **{** ⟨ *field-pat* ⟩ **}**]   test predicate abstraction *p*
specializer ::= *pred-spec* [**{** ⟨ *field-pat* ⟩ **}**]
pred-spec ::= *c*   *expr* **@** *c* is a class test
    | *p*   *expr* **@** *p***{** . . . **}** is alternate syntax
    for *p*(*expr*) **=>** **{** . . . **}**

A predicate abstraction takes a list of arguments and succeeds or fails as determined by its own predicate expression. A succeeding predicate abstraction invocation can return any value computed in its predicate expression, and the caller can retrieve any subset of the predicate abstraction's result bindings. Predicate signatures specify the type interface used in typechecking predicate abstraction callers and implementations. In this presentation, we prohibit recursive predicates.

Simple predicate abstractions are used just like ordinary classes:

```
predicate on_x_axis(p@point)
  when (p@cartesianPoint and test(p.y == 0))
    or (p@polarPoint and (test(p.theta == 0) or test(p.theta == pi)));

method draw(p@point) { ... } -- draw the point
method draw(p@on_x_axis) { ... } -- use a contrasting color so point is visible
```

In the following example, `CFG_2succ` is a CFG node with two successors. Each successor is marked with whether it is a loop exit (information which, in our implementation, is dynamically maintained when the CFG is modified) and the greatest loop it does not exit. It is advantageous for an iterative dataflow algorithm to propagate values along the loop exit only after reaching a fixed point within the loop; such an algorithm would dispatch on the `LoopExit` predicate. Similarly, the algorithm could switch from iterative to non-iterative mode when exiting the outermost loop, as indicated by `TopLevelLoopExit`.

```
predsignature LoopExit(CFGnode)
  return { loop:CFGloop };
predicate LoopExit(n@CFG_2succ{ next_true: t, next_false: f })
  when test(t.is_loop_exit) or (test_f.is_loop_exit)
```

```
    return { loop := outermost(t.containing_loop, f.containing_loop) };
  predicate TopLevelLoopExit(n@LoopExit{ loop@TopLevelScope });
```

2.5 Classifiers

Classifiers are a convenient syntax for imposing a linear ordering on a collection of predicates, ensuring mutual exclusion. They combine the state testing of predicate classes and the total ordering of pattern matching. An optional `otherwise` case, which executes if none of the predicates in the classifier evaluates to true, adds the guarantee of exhaustion. Multiple independent classifications of a particular class or object do not interfere with one another.

classifier-decl ::= **classify (⟨** *formal-pattern* **⟩)**
⟨ **as** p **when** *pred-expr* [**return {** ⟨ f := *expr* ⟩ **}**] ⟩
[**as** p **otherwise** [**return {** ⟨ f := *expr* ⟩ **}**]]

Here is an example of the use of classifiers:

```
class Window { ... }

classify(w@Window)
  as Iconified when test(w.iconified)
  as FullScreen when test(w.area() == RootWindow.area())
  as Big when test(w.area() > RootWindow.area()/2)
  as Small otherwise;

method move(w@FullScreen, x@int, y@int) { }          -- nothing to do
method move(w@Big, x@int, y@int) { ... }             -- move a wireframe outline
method move(w@Small, x@int, y@int) { ... }           -- move an opaque window
method move(w@Iconified, x@int, y@int) { ... }       -- modify icon, not window, coordinates

 -- resize, maximize, iconify similarly test these predicates
```

To force the classification to be mutually exclusive, each case is transformed into a predicate which includes the negation of the disjunction of all previous predicates. Therefore, an object is classified by some case only when it cannot be classified by any earlier case.

3 Dynamic and static semantics

The rest of this paper formalizes the dynamic and static semantics of a core predicate dispatching sublanguage. Figure 2 presents the abstract syntax of the core sublanguage. Appendix A defines desugaring rules that translate the high-level syntax of Figure 5 into the core syntax.

In the sequel, we assume that all variables are distinct so that the semantic rules can ignore the details of avoiding variable capture.

3.1 Dynamic semantics

This section explains how to select the most-specific applicable method at each message send. This selection relies on two key tests on predicated methods: whether a method is applicable to a call, and whether one method overrides another.

A method is applicable if its predicate evaluates to *true*; predicate evaluation also provides an extended environment in which the method's body is executed. Figure 3 defines the execution model of predicate evaluation in terms of the elaboration operator $\Rightarrow$ and several helper functions. We say $\langle P, K \rangle \Rightarrow \langle b, K' \rangle$ when the predicate P evaluates in the environment K to the boolean result b, producing the new environment K'. If the result b is *false*, then the resulting environment K' is ignored.

method-sig	::= **signature** m **(** $\langle\ T\ \rangle$ **)** **:** T	
method-decl	::= **method** m **(** $\langle\ v\ \rangle$ **)** **when** *pred-expr method-body*	
pred-expr	::= **true**	always applies
	$\mid$ **test** v	applies if v is *true*
	$\mid$ v **isa** c	applies if v is an instance of c or a subclass
	$\mid$ **let** v **:=** E	bind v to E; always applies
	$\mid$ p **(** $\langle\ v\ \rangle$ **)** **=>** **{** $\langle\ f = v\ \rangle$ **}**	test predicate abstraction p
	$\mid$ **not** *pred-expr*	negation
	$\mid$ *pred-expr* **and** *pred-expr*	conjunction (short-circuited)
	$\mid$ *pred-expr* **or** *pred-expr*	disjunction (short-circuited)
pred-sig	::= **predsignature** p **(** $\langle\ T\ \rangle$ **)** **return** **{** $\langle\ f : T\ \rangle$ **}**	
pred-decl	::= **predicate** p **(** $\langle\ v\ \rangle$ **)** **when** P **return** **{** $\langle\ f := v\ \rangle$ **}**	

Figure 2: Abstract syntax of the core language. Words and symbols in **boldface** represent terminals. Angle brackets denote zero or more comma-separated repetitions of an item. Square brackets contain optional expressions. We freely use parentheses around *pred-expr*s to indicate order of operations. Recursive predicates are forbidden.

Predicate dispatching considers one method m_1 to override another method m_2 exactly when m_1's predicate implies m_2's predicate and not vice versa. Section 4 describes how to compute the overriding relation, which can be performed at compile time.

Given the evaluation model for predicate expressions and the ability to compare predicate expressions for overriding, the execution of generic function invocations is straightforward. Suppose that generic function m is defined with the following cases:

$$\begin{array}{l}
\textbf{method}\, m(v_1, \ldots, v_n)\, \textbf{when}\, P_1\ \{B_1\} \\
\textbf{method}\, m(v_1, \ldots, v_n)\, \textbf{when}\, P_2\ \{B_2\} \\
\vdots \\
\textbf{method}\, m(v_1, \ldots, v_n)\, \textbf{when}\, P_k\ \{B_k\}
\end{array}$$

To evaluate the invocation $m(E_1, \ldots, E_n)$ in the environment K, we first obtain $\alpha_i = \mathrm{eval}(E_i, K)$ for all $i = 1, \ldots, n$. Then, for $j = 1, \ldots, k$, we obtain a truth value b_j and a new environment K_j through $\langle P_j, K[v_1 := \alpha_1, \ldots, v_n := \alpha_n]\rangle \Rightarrow \langle b_j, K_j\rangle$.[2]

Now let I be the set of integers i such that $b_i = true$, and find $i_0 \in I$ such that P_{i_0} overrides all others in $\{P_i\}_{i \in I}$. The result of evaluating $m(E_1, \ldots, E_n)$ is then the result of evaluating B_{i_0} in the environment K_{i_0}, so that variables bound in the predicate can be referred to in the body. If no such i_0 exists, then an exception is raised (a "message not understood" error if I is empty, or a "message ambiguous" error if there is no unique minimal element of I).

A clever implementation can make a number of improvements to this base algorithm. Here we briefly mention just a few such optimizations. First, common subexpression elimination over predicate expressions can limit the computation done in evaluating guards. Second, precomputed implication relationships can prevent the necessity for evaluating every predicate expression. If a more specific one is true, then the less specific one is certain to be satisfied; however, such satisfaction is irrelevant since the more specific predicate will be chosen. Third, clauses and methods can be reordered to succeed or fail more quickly.

3.2 Static semantics and typechecking

The operational model of predicate dispatch described in Section 3.1 can raise a run-time exception at a message send if no method is applicable or if no applicable method overrides all the others. We extend the typechecking rules of the underlying language to guarantee that no such exception occurs.

Figure 4 presents the static semantic domains, helper functions, and typechecking rules for the core predicate dispatching sublanguage.

---

[2] Since we assume that all variables are distinct, we can safely use the dynamic environment at the call site instead of preserving the static environment at the predicate abstraction's definition point.

$\alpha, \beta \in \mathit{value}$	Values in the underlying programming language
$b \in \{\mathit{true}, \mathit{false}\}$	Mathematical booleans
$K \in (\mathit{var\text{-}id} \rightarrow \mathit{value}) \cup (\mathit{pred\text{-}id} \rightarrow \mathit{pred\text{-}decl})$	Environments mapping variables to values and predicate names to predicate declarations

$\mathrm{lookup}(v, K) \rightarrow \alpha$	Look up the value of variable v in the environment K, returning the value α.
$K[v := \alpha] \rightarrow K'$	Bind the name v to the value α in the environment K, resulting in the new environment K'. If v is already bound, the existing binding is overridden.
$\mathrm{eval}(E, K) \rightarrow \alpha$	Evaluate the expression E in the environment K, returning the value α.
$\mathrm{instanceof}(\alpha, c) \rightarrow b$	Determine whether the value α is an instance of c or any subclass of c.
$\mathrm{accept}(\alpha) \rightarrow b$	Coerce arbitrary program values to *true* or *false*.

$$\frac{}{\langle \mathbf{true}, K\rangle \Rightarrow \langle \mathit{true}, K\rangle}$$

$$\frac{\mathrm{lookup}(v, K) = \alpha \qquad \mathrm{accept}(\alpha) = b}{\langle v, K\rangle \Rightarrow \langle b, K\rangle}$$

$$\frac{\mathrm{lookup}(v, K) = \alpha \qquad \mathrm{instanceof}(\alpha, c) = b}{\langle v\ \mathbf{isa}\ c, K\rangle \Rightarrow \langle b, K\rangle}$$

$$\frac{\mathrm{eval}(E, K) = \alpha \qquad K[v := \alpha] = K'}{\langle \mathbf{let}\ \ v\ \mathbf{:=}\ E, K\rangle \Rightarrow \langle \mathit{true}, K'\rangle}$$

$$\frac{\begin{array}{c} \forall i \in \{1, \ldots, n\} \quad \mathrm{eval}(v'_i, K) = \alpha_i \\ \mathrm{lookup}(p, K) = \mathbf{predicate}\, p(v_1, \ldots, v_n)\ \mathbf{when}\ \ P\, \mathbf{return}\ \ \mathbf{\{} f_1\ \mathbf{:=}\ E_1, \ldots, f_m\ \mathbf{:=}\ E_m, \ldots \mathbf{\}} \\ \langle P, K[v_1 := \alpha_1, \ldots, v_n := \alpha_n]\rangle \Rightarrow \langle \mathit{false}, K'\rangle \end{array}}{\langle p(v'_1, \ldots, v'_n)\ \mathbf{=>}\ \mathbf{\{} f_1 = w_1, \ldots, f_m = w_m \mathbf{\}}, K\rangle \Rightarrow \langle \mathit{false}, K\rangle}$$

$$\frac{\begin{array}{c} \forall i \in \{1, \ldots, n\} \quad \mathrm{eval}(v'_i, K) = \alpha_i \\ \mathrm{lookup}(p, K) = \mathbf{predicate}\, p(v_1, \ldots, v_n)\ \mathbf{when}\ \ P\, \mathbf{return}\ \ \mathbf{\{} f_1\ \mathbf{:=}\ E_1, \ldots, f_m\ \mathbf{:=}\ E_m, \ldots \mathbf{\}} \\ \langle P, K[v_1 := \alpha_1, \ldots, v_n := \alpha_n]\ \rangle \Rightarrow \langle \mathit{true}, K'\rangle \\ \forall i \in \{1, \ldots, m\} \quad \mathrm{eval}(E_i, K') = \beta_i \\ K[w_1 := \beta_1, \ldots, w_m := \beta_m] = K'' \end{array}}{\langle p(v'_1, \ldots, v'_n)\ \mathbf{=>}\ \mathbf{\{} f_1 = w_1, \ldots, f_m = w_m \mathbf{\}}, K\rangle \Rightarrow \langle \mathit{true}, K''\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle b, K'\rangle}{\langle \mathbf{not}\ P, K\rangle \Rightarrow \langle \neg b, K\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{false}, K'\rangle}{\langle P\ \mathbf{and}\ Q, K\rangle \Rightarrow \langle \mathit{false}, K\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{true}, K'\rangle \qquad \langle Q, K'\rangle \Rightarrow \langle \mathit{false}, K''\rangle}{\langle P\ \mathbf{and}\ Q, K\rangle \Rightarrow \langle \mathit{false}, K\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{true}, K'\rangle \qquad \langle Q, K'\rangle \Rightarrow \langle \mathit{true}, K''\rangle}{\langle P\ \mathbf{and}\ Q, K\rangle \Rightarrow \langle \mathit{true}, K''\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{true}, K'\rangle}{\langle P\ \mathbf{or}\ Q, K\rangle \Rightarrow \langle \mathit{true}, K\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{false}, K'\rangle \qquad \langle Q, K\rangle \Rightarrow \langle \mathit{true}, K''\rangle}{\langle P\ \mathbf{or}\ Q, K\rangle \Rightarrow \langle \mathit{true}, K\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{false}, K'\rangle \qquad \langle Q, K\rangle \Rightarrow \langle \mathit{false}, K''\rangle}{\langle P\ \mathbf{or}\ Q, K\rangle \Rightarrow \langle \mathit{false}, K\rangle}$$

Figure 3: Dynamic semantic domains, helper functions, and evaluation rules

$T \leq T'$	Type T is a subtype of T'.
conformant-type(T, c)	Return the most-specific type T' such that every subclass c' of c that conforms to T also conforms to T'. This helper function is supplied by the underlying programming language.
$, + , ' = , ''$	Overriding extension of typing environments. For each $v \in \mathrm{dom}(, ')$, if $, ' \models v : T'$, then $, '' \models v : T'$; for each $v \in \mathrm{dom}(,) \setminus \mathrm{dom}(Gamma')$, if $, \models v : T$, then $, '' \models v : T$.

$$\frac{}{, \vdash \textbf{signature}\ m(T_1, \ldots, T_n) \mathbf{:} T_r \Rightarrow , + \{m : (T_1, \ldots, T_n) \rightarrow T_r\}}$$

$$\frac{\begin{array}{c} , \models m : (T_1, \ldots, T_n) \rightarrow T_r \\ , + \{v_1 : T_1, \ldots, v_n : T_n\} \vdash P \Rightarrow , ' \qquad , ' \models \textit{method-body} : T_b \qquad T_b \leq T_r \end{array}}{, \vdash \textbf{method}\ m(v_1, \ldots, v_n)\ \ \textbf{when}\ \ P\ \textit{method-body} \Rightarrow ,}$$

$$\frac{}{, \vdash \textbf{predsignature}\ p(T_1, \ldots, T_n)\ \ \textbf{return}\ \mathbf{\{} f_1 \mathbf{:} T_1^r, \ldots, f_m \mathbf{:} T_m^r \mathbf{\}} \Rightarrow , + \{p : (T_1, \ldots, T_n) \rightarrow \{f_1 : T_1^r, \ldots, f_m : T_m^r\}\ \}}$$

$$\frac{\begin{array}{c} , \models p : (T_1, \ldots, T_n) \rightarrow \{f_1 : T_1^r, \ldots, f_m : T_m^r, \ldots\} \\ , + \{v_1 : T_1, \ldots, v_n : T_n\} \vdash P \Rightarrow , ' \\ \forall i \in \{1, \ldots, m\} \quad , ' \models v_i' : T_i' \ \wedge\ T_i' \leq T_i^r \end{array}}{, \vdash \textbf{predicate}\ p(v_1, \ldots, v_n)\ \ \textbf{when}\ \ P\ \ \textbf{return}\ \mathbf{\{} f_1 \mathbf{:=} v_1', \ldots, f_m \mathbf{:=} v_m' \mathbf{\}} \Rightarrow ,}$$

$$\frac{}{, \vdash \textbf{true} \Rightarrow ,}$$

$$\frac{, \models v : \mathit{Bool}}{, \vdash v \Rightarrow ,}$$

$$\frac{, \models v : T \qquad \textit{conformant-type}(c, T) = T'}{, \vdash v\ \textbf{isa}\ c \Rightarrow , + \{v : T'\}}$$

$$\frac{, \models \mathit{expr} : T}{, \vdash \textbf{let}\ \ v \mathbf{:=} \mathit{expr} \Rightarrow , + \{v : T\}}$$

$$\frac{\begin{array}{c} , \models p : (T_1, \ldots, T_n) \rightarrow \{f_1 : T_1^r, \ldots, f_m : T_m^r, \ldots\} \\ , \models v_1 : T_1' \quad \ldots \quad , \models v_n : T_n' \qquad T_1' \leq T_1 \quad \ldots \quad T_n' \leq T_n \end{array}}{, \vdash p(v_1, \ldots, v_n) \mathbf{=>} \mathbf{\{} f_1 = v_1', \ldots, f_m = v_m' \mathbf{\}} \Rightarrow , + \{v_1' : T_1^r, \ldots, v_m' : T_m^r\}}$$

$$\frac{, \vdash P \Rightarrow , '}{, \vdash \textbf{not}\ P \Rightarrow ,}$$

$$\frac{, \vdash P_1 \Rightarrow , ' \qquad , ' \vdash P_2 \Rightarrow , ''}{, \vdash P_1\ \textbf{and}\ P_2 \Rightarrow , ''}$$

$$\frac{, \vdash P_1 \Rightarrow , ' \qquad , \vdash P_2 \Rightarrow , ''}{, \vdash P_1\ \textbf{or}\ P_2 \Rightarrow ,}$$

Figure 4: Typechecking rules. The hypothesis $, \models \mathit{expr} : T$ indicates that typechecking in typing environment , assigns type T to *expr*. The judgment $, \vdash P \Rightarrow , '$ represents extension of typechecking environments: given type environment , , P typechecks and produces new typechecking environment $, '$. The return type for a predicate invocation is an unordered record.

We can separate typechecking into two parts: *client-side*, which handles all checking of expressions in the underlying language and uses method signatures to typecheck message sends, and *implementation-side*, which checks method and predicate implementations against their corresponding signatures. Only implementation-side checking is affected by predicate dispatching.

Implementation-side typechecking must guarantee *completeness* and *uniqueness*. Completeness guarantees that no "message not understood" error is raised: for every possible set of arguments at each call site, some method is applicable. Let P_m be the disjunction of the predicates of all of m's implementations, and let P_s be a predicate expressing the set of argument classes that conform to the types in the method signature. (See below for the details of predicate P_s; a class c conforms to a type T if every object which is an instance of that class has type T or a subtype of T.) If P_s implies P_m, then some method is always applicable. Uniqueness guarantees that no "message ambiguous" error is raised: for no possible set of arguments at any call site are there multiple most-specific methods. Uniqueness is guaranteed if, for each pair of predicates P and Q attached to two different implementations, either P and Q are disjoint (so their associated methods can never be simultaneously applicable) or one of the predicates implies the other (so one of the methods overrides the other). Section 4 presents implication and disjointness tests over predicate expressions.

Completeness checking requires a predicate P_s that expresses the set of tuples of values $v_1, \ldots, v_n$ conforming to some signature's argument types $T_1, \ldots, T_n$; this predicate depends on the host language's model of classes and typing. If classes and types are the same, and all classes are concrete, then the corresponding predicate is simply v_1 **isa** T_1 **and** ... **and** v_n **isa** T_n. If abstract classes are allowed, then each v_i **isa** T_i is replaced with v_i **isa** T_{i1} **or** ... **or** v_i **isa** T_{im}, where the T_{ij} are the top concrete subclasses of T_i. If inheritance and subtyping are separate notions, then the predicates become more complex.

Our typechecking need not test that methods conform to signatures, unlike previous work on typechecking multimethods [CL95]. In predicate dispatching, a method's formal argument has two distinct types: the "external" type derived from the signature declaration, and the possibly finer "internal" type guaranteed by successful evaluation of the method's predicate. The individual **isa** tests narrow the type of the tested value to the most-specific type to which all classes passing the test conform, in a host-language-specific manner, using *conformant-type*. The *conformant-type* function replaces the more complicated conformance test of earlier work.

4 Comparing predicate expressions

The static and dynamic semantics of predicate dispatching require compile-time tests of implication between predicates to determine the method overriding relationship. The static semantics also requires tests of completeness and uniqueness to ensure the absence of message-not-understood errors and message-ambiguous errors, respectively. All of these tests reduce to tautology tests over predicates. Method m_1 overrides method m_2 iff m_1's predicate implies that of m_2 — that is, if (**not** m_1) **or** m_2 is true. A set of methods is complete if the disjunction of their predicates is true. Uniqueness for a set of methods requires that for any pair of methods, either one overrides the other, or the two are logically exclusive. Two formulas are mutually exclusive exactly if one implies the negation of the other.

Section 4.1 presents a simple, sound, complete tautology test over predicate expressions. Because determining logical tautology is NP-complete, in the worst case an algorithm takes exponential time in the size of the predicate expressions. For object-oriented dispatch, this is the number of arguments to a method (a small constant). Simple optimizations (Section 4.2) make the tests fast in many practical situations. This cost is incurred only at compile time; at run time, precomputed overriding relations among methods are simply looked up.

We treat expressions from the underlying programming language as black boxes (but do identify those which perform the same computation). Tests involving the run-time values of arbitrary host language expressions are undecidable. The algorithm presented here also does not address recursive predicates. While we have a set of heuristics that succeed in many common, practical cases, we do not yet have a complete, sound, efficient algorithm.

4.1 The base algorithm

The base algorithm for testing predicate tautology has three components. First, the predicate expression is canonicalized to macro-expand predicate abstractions, eliminate variable bindings, and use canonical names for formal arguments. This transformation prevents different names for the same value from being considered distinct. Second, implication relations are computed among the atomic predicates (for instance, x **isa** `int` implies x **isa** `num`). Finally, the canonicalized predicate is tested for every assignment of atomic predicates to truth values which is consistent with the atomic predicate implications. The predicate is a tautology iff evaluating it in every consistent truth assignment yields *true*.

4.1.1 Canonicalization

Canonicalization performs the following transformations:

- Expand predicate calls inline, replacing the => clause by a series of **let** bindings.
- Replace **let**-bound variables by the expressions to which they are bound, and replace **let** expressions by **true**.
- Canonically rename formal parameters according to their position in the formal list.

After canonicalization, each predicate expression is a logical formula over the following atoms with connectives **and**, **or**, and **not**.

$$\begin{array}{lcl} \textit{pred-atom} & ::= & \textbf{true} \\ & | & \textbf{test}\ E \\ & | & E\ \textbf{isa}\ c \end{array}$$

Canonicalized predicates are a compile-time construct used only for predicate comparison; they are never executed. Canonicalized predicates bind no variables, and they use only global variables and formal parameters.

In the worst case, canonicalization exponentially blows up expression sizes. For instance, in

$$\textbf{let}\ x_1 = x + x\ \textbf{and let}\ x_2 = x_1 + x_1\ \textbf{and let}\ x_3 = x_2 + x_2\ \textbf{and}\ \ldots\ \textbf{and test}\ \ x_n = y \quad ,$$

the final x_n is replaced by an expression containing 2^n instances of x. Inline expansion of predicate abstractions similarly contributes to this blowup. As with ML typechecking [KM89], which is exponential in the worst case but linear in practice, we anticipate that predicates leading to exponential behavior will be rare.

In the sequel we will consider two expressions identical if, after canonicalization, they have the same abstract syntax tree.

Omitting this step prevents some equivalent expressions from being recognized as such, but does not prevent the remainder of the algorithm from succeeding when results are named and reused rather than the computation repeated.

4.1.2 Truth assignment checking

This section presents a simple exponential-time algorithm to check logical tautology; because the problem is NP-complete, any algorithm takes exponential time in the worst case. Let there be n distinct predicate atoms in the predicate; there are 2^n different truth assignments for those atoms. Not all of those truth assignments are consistent with the implications over predicate atoms: for instance, it is not sensible to set `a` **isa** `int` to *true* but `a` **isa** `num` to *false*, because `a` **isa** `int` implies `a` **isa** `num`. If every consistent truth assignment satisfies the predicate, then the predicate is a tautology. Each check of a single truth assignment takes time linear in the size of the predicate expressions, for a total time of $O(n2^n)$.

The following rules specify implication over (possibly negated) canonical predicate atoms.

- $E_1\ \textbf{isa}\ c_1 \Rightarrow E_2\ \textbf{isa}\ c_2$ iff ($E_1 \equiv E_2$) and (c_1 is a subclass of c_2)
- $E_1\ \textbf{isa}\ c_1 \Rightarrow \textbf{not}(E_2\ \textbf{isa}\ c_2)$ iff ($E_1 \equiv E_2$) and (c_1 is disjoint from c_2)
- $a_1 \Rightarrow a_2$ iff $\textbf{not}\ a_2 \Rightarrow \textbf{not}\ a_1$
- $a_1 \Rightarrow \textbf{not}\ a_2$ iff $a_2 \Rightarrow \textbf{not}\ a_1$

Two classes are disjoint if they have no common descendant. As is usual, $\textbf{not}\,\textbf{not}\ a = a$.

4.2 Optimizations

The worst-case exponential-time cost to check predicate tautology need not prevent its use in practice. Satisfiability is checked only at compile time. When computing overriding relationships, the predicates tend to be small (linear in the number of arguments to a method). We present heuristics that reduce the costs even further.

Logical simplification — such as eliminating uses of **true**, **not not** a, and a **and not** a — can be performed as part of canonicalization to reduce the size of predicate expressions.

Unrelated atomic predicates can be treated separately. To determine whether **method** $m_1(f_1@c_1, f_2@c_2)\{...\}$ overrides **method** $m_1(f_1@c_3, f_2@c_4)\{...\}$ it is sufficient to independently determine the relationship between c_1 and c_3 and that between c_2 and c_4. Two tests with a smaller exponent replace one with a larger one, substantially reducing the overall cost. This technique always solves ordinary single and multiple dispatching overriding in time constant and linear in the number of formals, respectively, by examining each formal position independently. The technique also applies to more complicated cases, by examining subsets of formal parameters which appear together in tests from the underlying programming language.

It is not always necessary to completely expand predicate abstraction calls as part of canonicalization. If relations between predicate abstractions or other predicate expressions are known, then the tautology test can use them directly. As one example, no complicated test is required in order to determine that different cases of a classifier are mutually exclusive, as that property is satisfied by definition.

The side conditions on atomic predicate values (their implication relationships) usually prevent the need to check all 2^n different truth assignments for a predicate containing n atomic predicates. When `a` **`isa`** `int` is set to *true*, then all truth assignments which set `a` **`isa`** `num` to *false* can be skipped without further consideration.

Finally, it may be possible to achieve faster results in some cases by recasting the tautology test. Rather than attempting to prove that every truth assignment satisfies a predicate expression, it may be advantageous to search for a single truth assignment that satisfies its negation.

5 Related work

5.1 Object-oriented approaches

In the model of predicate dispatching, traditional object-oriented dispatching translates to either a single class test on the receiver argument or, for multiple dispatching, a conjunction of class tests over several arguments. Full predicate dispatching additionally enables testing arbitrary boolean expressions from the underlying language; accessing and naming subcomponents of the arguments; performing tests over multiple arguments; and arbitrarily combining tests via conjunction, disjunction, and negation. Also, named predicate abstractions effectively introduce new virtual classes and corresponding subclassing links into the program inheritance hierarchy. Predicate dispatching preserves the ability in object-oriented languages to statically determine when one method overrides another and when no message lookup errors can occur. Singly-dispatched object-oriented languages have efficient method lookup algorithms and separate typechecking, which depend crucially on the absence of any separate modules that dispatch on other argument positions. Multiply-dispatched object-oriented languages have more challenging problems in implementation [KR89, CTK94, AGS94] and typechecking [CL95], and predicate dispatching in its unrestricted form shares these challenges.

Predicate classes [Cha93b] are an earlier extension of object-oriented dispatching to include arbitrary boolean predicates. A predicate class which inherits from some class A and has an associated predicate expression *guard* would be modeled as a named predicate abstraction that tests `@`A `and` *guard*. Predicate dispatching is more general, for example by being able to define predicates over multiple arguments. Predicate dispatching exploits the structure of **and**, **or**, and **not** to automatically determine when no message lookup errors can occur, while predicate classes rely on uncheckable user assertions about the relations between the predicate classes' guard expressions in order to do typechecking.

Classifiers in Kea [HHM90b, HHM90a, MHH91] let an instance of a class be dynamically reclassified as being of a subclass. A classifier for a class is composed of a sequence of arbitrary-predicate/subclass pairs, with an object of the input class automatically classified as being of the subclass with the first successful predicate. Because the sequence of predicates is totally ordered and the first successful predicate takes precedence over all later predicates, a classifier provides a concise syntax for a set of mutually exclusive, exhaustive predicate abstractions. Predicate abstractions are more general than classifiers in many of the ways discussed above, but they also provide syntactic support for this important idiom. Kea is a purely functional language, so classifiers do not need to consider the semantics of reclassifying objects when the values of predicates change; predicate dispatching addresses this issue by (conceptually) performing reclassification as needed as part of message dispatching.

Modes [Tai93] are another mechanism for adding dynamic reclassification of a class into a subclass. Unlike predicate classes and classifiers, the modes of a class are not first-class subclasses but rather internal components of a class that cannot be extended externally and that cannot exploit inheritance to factor shared code. Mode reselection can be done either explicitly at the end of each method or implicitly after each assignment using a declaratively specified classification.

5.2 Pattern matching approaches

Predicate dispatching supports many of the facilities found in pattern matching as in ML [MTH90] and Haskell [HJW+92], including tests over arbitrary nested structure, binding of names to subcomponents, and arbitrary boolean guard expressions. Predicate dispatching additionally supports inheritance (its class tests are more general than datatype constructor patterns), disjunctions and negations of tests and conjunctions of tests on the same object, and named predicate abstractions to factor out common patterns of tests and to offer conditional views of objects extended with virtual fields. The patterns in a function are totally ordered, while predicate dispatching computes a partial order over predicates and warns when two patterns might be ambiguous. Finally, new methods can be added to existing generic functions without changing any existing code, while new patterns can be added to a function only by modifying it.

Views [Wad87] extend pattern matching to abstract data types by allowing an abstract data type to offer a number of views of itself as a concrete datatype, over which pattern matching is defined. Predicate dispatching supports "pattern matching" over the results of methods (by **let**-binding their results to names and then testing those names, just as field contents are bound and tested), and those methods can serve as accessor functions to a virtual view of the object, for instance `rho` and `theta` methods presenting a polar view of a cartesian point. Views must be isomorphisms, which enables equational reasoning over them; by contrast, named predicate abstractions provide conditional views of an object without requiring the presence of both in and out views.

Pizza [OW97] supports both algebraic datatypes (and associated pattern matching) and object-oriented dispatching, but the two mechanisms are largely distinct. The authors argue that datatypes are good for fixed numbers of representations with extensible operations, while classes are good for a fixed set of operations with extensible representations. By integrating pattern matching and dispatching, including multimethods, predicate dispatching achieves extensibility in both dimensions along with the syntactic convenience of pattern matching. Predicate dispatching faces more difficult implementation and separate typechecking challenges with the shift to multimethod-like dispatching.

6 Conclusions

Many language features express the concept of selecting a most-specific applicable method from a collection of candidates, including object-oriented dispatch, pattern matching, views, predicate classes, and classifiers. Predicate dispatching integrates and generalizes these mechanisms in a single framework, based on a core language of boolean expressions over class tests and arbitrary expressions, explicit binding forms to generalize features of pattern matching, and named predicate abstractions with result bindings. By providing a single integrated mechanism, programs can then take advantage of various styles of dispatch and even combine them to create applicability conditions that were previously impossible or inconvenient to express.

We have implemented predicate dispatching in the context of Dubious, a simple core multiply-dispatched object-oriented programming language. The implementation supports all the examples presented in this paper, though for clarity this paper uses a slightly different presentation syntax. The implementation includes the complete, sound satisfiability test of Section 4 and some of the optimizations of Section 3.1, but few optimizations from Section 4.2. This implementation was helpful in verifying our base design. We expect that it will also provide insight into the advantages and disadvantages of programming with predicate dispatching, as well as help us to evaluate optimization strategies.

So far, we have focused on developing the static and dynamic semantics for predicate dispatching. Two unresolved practical issues that we will address in the future are efficient implementation techniques and separate typechecking support for predicate dispatching. We anticipate that efficient implementations of unrestricted predicate dispatching will build upon work on efficient implementation of multimethod dispatching and on predicate classes. In addition, static analyses that factor a collection of predicates to avoid redundant tests and side-effect analyses that determine when predicates need not be re-evaluated appear to be promising lines for future research. Similarly, separate typechecking of collections of predicated methods will build upon current work to develop modular and incremental methods for typechecking multimethods [CL95].

Acknowledgments

We thank Todd Millstein, Vassily Litvinov, and Wilson Hsieh for their comments on a draft of this paper. This research is supported in part by an NSF grant (number CCR-9503741), an NSF Young Investigator Award (number CCR-9457767), a grant from the Office of Naval Research (contract number N00014-94-1-1136), and gifts from Sun Microsystems, IBM, Xerox PARC, Pure Software, and Edison Design Group.

References

[AGS94] Eric Amiel, Olivier Gruber, and Eric Simon. Optimizing multi-method dispatch using compressed dispatch tables. In *Proceedings OOPSLA '94*, pages 244–258, Portland, OR, October 1994.

[BKK+86] Daniel G. Bobrow, Ken Kahn, Gregor Kiczales, Larry Masinter, Mark Stefik, and Frank Zdybel. Commonloops: Merging lisp and object-oriented programming. In *Proceedings OOPSLA '86*, pages 17–29, November 1986. Published as ACM SIGPLAN Notices, volume 21, number 11.

[Cha92] Craig Chambers. Object-oriented multi-methods in Cecil. In O. Lehrmann Madsen, editor, *Proceedings ECOOP '92*, LNCS 615, pages 33–56, Utrecht, The Netherlands, June 1992. Springer-Verlag.

[Cha93a] Craig Chambers. The Cecil language: Specification and rationale. Technical Report TR-93-03-05, Department of Computer Science and Engineering. University of Washington, March 1993.

[Cha93b] Craig Chambers. Predicate classes. In O. Nierstrasz, editor, *Proceedings ECOOP '93*, LNCS 707, pages 268–296, Kaiserslautern, Germany, July 1993. Springer-Verlag.

[CL95] Craig Chambers and Gary T. Leavens. Typechecking and modules for multi-methods. *ACM Transactions on Programming Languages and Systems*, 17(6):805–843, November 1995.

[CTK94] Weimin Chen, Volker Turau, and Wolfgang Klas. Efficient dynamic look-up strategy for multi-methods. In M. Tokoro and R. Pareschi, editors, *Proceedings ECOOP '94*, LNCS 821, pages 408–431, Bologna, Italy, July 1994. Springer-Verlag.

[HHM90a] J. Hamer, J.G. Hosking, and W.B. Mugridge. A method for integrating classification within an object-oriented environment. Technical Report Auckland Computer Science Report No. 48, Department of Computer Science, University of Auckland, October 1990.

[HHM90b] J.G. Hosking, J. Hamer, and W.B. Mugridge. Integrating functional and object-oriented programming. In *Technology of Object-Oriented Languages and Systems TOOLS 3*, pages 345–355, Sydney, 1990.

[HJW+92] Paul Hudak, Simon Peyton Jones, Philip Wadler, Brian Boutel, Jon Fairbairn, Joseph Fasel, Maria Guzman, Kevin Hammond, John Hughes, Thomas Johnsson, Dick Kieburtz, Rishiyur Nikhil, Will Partain, and John Peterson. Report on the programming language Haskell, version 1.2. *SIGPLAN Notices*, 27(5), May 1992.

[KM89] Paris C. Kanellakis and John C. Mitchell. Polymorphic unification and ML typing. In ACM-SIGPLAN ACM-SIGACT, editor, *Conference Record of the 16th Annual ACM Symposium on Principles of Programming Languages (POPL '89)*, pages 105–115, Austin, TX, USA, January 1989. ACM Press.

[KR89] Gregor Kiczales and Luis Rodriguez. Efficient method dispatch in PCL. Technical Report SSL 89-95, Xerox PARC Systems Sciences Laboratory, 1989.

[MHH91] Warwick B. Mugridge, John Hamer, and John G. Hosking. Multi-methods in a statically-typed programming language. In P. America, editor, *Proceedings ECOOP '91*, LNCS 512, pages 307–324, Geneva, Switzerland, July 15-19 1991. Springer-Verlag.

<table>
<tr><td><i>method-sig</i></td><td>::=</td><td>signature <i>m</i> (⟨ <i>T</i> ⟩) : <i>T</i></td><td></td></tr>
<tr><td><i>method-decl</i></td><td>::=</td><td>method <i>m</i> (⟨ <i>formal-pattern</i> ⟩) [when <i>pred-expr</i>] <i>method-body</i></td><td></td></tr>
<tr><td><i>pred-sig</i></td><td>::=</td><td>predsignature <i>p</i> (⟨ <i>T</i> ⟩) return { ⟨ <i>f</i> : <i>T</i> ⟩ }</td><td></td></tr>
<tr><td><i>pred-decl</i></td><td>::=</td><td>predicate <i>p</i> (⟨ <i>formal-pattern</i> ⟩)
[when <i>pred-expr</i>] [return { ⟨ <i>f</i> := <i>expr</i> ⟩ }]</td><td></td></tr>
<tr><td><i>classifier-decl</i></td><td>::=</td><td>classify (⟨ <i>formal-pattern</i> ⟩)
⟨ as <i>p</i> when <i>pred-expr</i> [return { ⟨ <i>f</i> := <i>expr</i> ⟩ }] ⟩
[as <i>p</i> otherwise [return { ⟨ <i>f</i> := <i>expr</i> ⟩ }]]</td><td></td></tr>
<tr><td><i>pred-expr</i></td><td>::=</td><td>true</td><td>always succeeds</td></tr>
<tr><td></td><td>|</td><td>false</td><td>never succeeds</td></tr>
<tr><td></td><td>|</td><td><i>expr</i> @ <i>specializer</i></td><td>succeeds if <i>expr</i> evaluates to an instance or subclass of the specified class or predicate</td></tr>
<tr><td></td><td>|</td><td>test <i>expr</i></td><td>succeeds if <i>expr</i> evaluates to <i>true</i></td></tr>
<tr><td></td><td>|</td><td>let v := E</td><td>bind v to E; always succeeds</td></tr>
<tr><td></td><td>|</td><td><i>p</i> (⟨ <i>expr</i> ⟩) [=> { ⟨ <i>field-pat</i> ⟩ }]</td><td>test predicate abstraction <i>p</i></td></tr>
<tr><td></td><td>|</td><td>not <i>pred-expr</i></td><td>negation</td></tr>
<tr><td></td><td>|</td><td><i>pred-expr</i> and <i>pred-expr</i></td><td>conjunction (short-circuited)</td></tr>
<tr><td></td><td>|</td><td><i>pred-expr</i> or <i>pred-expr</i></td><td>disjunction (short-circuited)</td></tr>
<tr><td><i>formal-pattern</i></td><td>::=</td><td>[<i>v</i>] [@ <i>specializer</i>]</td><td>like <i>v</i> isa <i>specializer</i> in <i>pred-expr</i></td></tr>
<tr><td><i>field-pat</i></td><td>::=</td><td><i>m</i> [= <i>v</i>] [@ <i>specializer</i>]</td><td></td></tr>
<tr><td><i>specializer</i></td><td>::=</td><td><i>pred-spec</i> [{ ⟨ <i>field-pat</i> ⟩ }]</td><td></td></tr>
<tr><td><i>pred-spec</i></td><td>::=</td><td><i>c</i></td><td><i>expr</i> @ <i>c</i> is a class test</td></tr>
<tr><td></td><td>|</td><td><i>p</i></td><td><i>expr</i> @ <i>p</i>{ ... } is alternate syntax for <i>p</i>(<i>expr</i>) => { ... }</td></tr>
<tr><td></td><td>|</td><td>not <i>pred-spec</i></td><td>succeeds if <i>pred-expr</i> does not</td></tr>
<tr><td></td><td>|</td><td><i>pred-spec</i> & <i>pred-spec</i></td><td>succeeds if both <i>pred-expr</i>s do</td></tr>
<tr><td></td><td>|</td><td><i>pred-spec</i> | <i>pred-spec</i></td><td>succeeds if either <i>pred-expr</i> does</td></tr>
</table>

Figure 5: Full extended syntax for predicate dispatching. The notation is as for Figure 2. The syntax is as presented incrementally in Section 2, with a few additions (such as boolean operators in *pred-specs*).

[MTH90] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML.* MIT Press, 1990.

[OW97] Martin Odersky and Philip Wadler. Pizza into Java: Translating theory into practice. In *Conference Record of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 146–159, January 1997.

[Ste90] Guy L. Steele Jr. *Common Lisp: The Language.* Digital Press, Bedford, MA, 1990. Second edition.

[Tai93] Antero Taivalsaari. Object-oriented programming with modes. *Journal of Object-Oriented Programming*, pages 25–32, June 1993.

[Wad87] Philip Wadler. Views: A way for pattern matching to cohabit with data abstraction. In *Proceedings of the Fourteenth Annual ACM Symposium on Principles of Programming Languages*, pages 307–313, Munich, Germany, January 1987.

A Desugaring rules

The following rewrite rules desugar the high-level syntax of Figure 5 into the core abstract syntax of Figure 2. For brevity, we use $\bigwedge_{i=1}^{n} \{P_i\}$ to stand for the conjunction of the terms: P_1 **and**...**and** P_n. Variables v' and v'_i are new variables which do not appear elsewhere in the program. Ceiling braces $\lceil\cdot\rceil$ surround (potentially) sugared expressions; application of the rewrite rules eliminates those braces.

For brevity, we omit the rewrite rules which introduce defaults for omitted optional program fragments: dummy variables for pattern variables, "**@***any*" specializers, empty field pattern sets in specializers, and "**when true**" and "**return { }**" clauses. Additional rules may be introduced to simplify the resulting formula, such as converting "*v* **isa** *any*" to "**true**" and performing logical simplification.

Declarations: move specializers into **when** clause

$\lceil\textbf{method}\ m(v_1\ @\ S_1,\ldots,v_n\ @\ S_n)\ \ \textbf{when}\ \ P\ \ \textit{method-body}\rceil$

$\quad\Longrightarrow\ \ \textbf{method}\ m(v_1,\ldots,v_n)\ \ \textbf{when}\ \ \bigwedge_{i=1}^{n}\{\lceil v_i\ \textbf{isa}\ S_i\rceil\}\ \textbf{and}\ \lceil P\rceil\ \ \textit{method-body}$

$\lceil\textbf{predicate}\ p(v_1\ @\ S_1,\ldots,v_n\ @\ S_n)\ \ \textbf{when}\ \ P\ \ \textbf{return}\ \ \{f_1 := E_1,\ldots,f_m := E_m\}\rceil$

$\quad\Longrightarrow\ \ \textbf{predicate}\ p(v_1,\ldots,v_n)\ \ \textbf{when}\ \ \bigwedge_{i=1}^{n}\{\lceil v_i\ \textbf{isa}\ S_i\rceil\}\ \textbf{and}\ \ \lceil P\rceil\ \textbf{and}\ \ \lceil\textbf{return}\ \{f_1 := E_1,\ldots,f_m := E_m\}\rceil$

Logic

$$
\begin{aligned}
\lceil P_1\ \textbf{and}\ P_2\rceil &\ \Longrightarrow\ \lceil P_1\rceil\ \textbf{and}\ \lceil P_2\rceil\\
\lceil P_1\ \textbf{or}\ P_2\rceil &\ \Longrightarrow\ \lceil P_1\rceil\ \textbf{or}\ \lceil P_2\rceil\\
\lceil\textbf{not}\ P\rceil &\ \Longrightarrow\ \textbf{not}\ \lceil P\rceil\\
\lceil\textbf{false}\rceil &\ \Longrightarrow\ \textbf{not true}
\end{aligned}
$$

Name non-variable expressions

$$
\begin{aligned}
\lceil E\rceil &\ \Longrightarrow\ \textbf{let}\ \ v' := E\ \textbf{and}\ \ v'\\
\lceil E\ \textbf{isa}\ S\rceil &\ \Longrightarrow\ \textbf{let}\ \ v' := E\ \textbf{and}\ \lceil v'\ \textbf{isa}\ S\rceil\\
\lceil p(E_1,\ldots,E_n)\ \ \mathtt{=>}\ \ \{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil &\ \Longrightarrow\ \bigwedge_{i=1}^{n}\{\textbf{let}\ \ v_i' := E_i\}\\
&\qquad\ \ \textbf{and}\ \lceil p(v_1',\ldots,v_n')\ \ \mathtt{=>}\ \ \{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil\\
\lceil\textbf{return}\ \{f_1 := E_1,\ldots,f_m := E_m\}\rceil &\ \Longrightarrow\ \textbf{and}\bigwedge_{i=1}^{m}\{\textbf{let}\ \ v_i' := E_i\}\ \textbf{return}\ \{f_1 := v_1',\ldots,f_m := v_m'\}
\end{aligned}
$$

Field bindings: A field name is *generated* by Pname if Pname is a class containing the field, if Pname is a predicate name whose result contains the field, if Pname is a disjunction both of whose disjuncts generate the field, if Pname is a conjunction either of whose conjuncts generates the field, or if the field name is actually a single-argument method name.

$\lceil v\ \textbf{isa}\ c\ \ \{f_1 = v_1\ @\ S_1,\ldots,f_n = v_n\ @\ S_n\}\rceil$

$\quad\Longrightarrow\ \ v\ \textbf{isa}\ c\ \textbf{and}\ \ \bigwedge_{i=1}^{n}\{\lceil v.f_i\ \textbf{isa}\ S_i\rceil\}$

$\lceil v\ \textbf{isa}\ P\ \ \{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m,\ldots,\mathit{Fdpat}_n\}\rceil\qquad \mathit{Fdpat}_i$ is generated by $P \neq c$ for $1 \leq i \leq m < n$

$\quad\Longrightarrow\ \ \lceil v\ \textbf{isa}\ P\ \ \{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil\ \textbf{and}\ \lceil v\ \textbf{isa}\ \mathit{any}\ \ \{\mathit{Fdpat}_{m+1},\ldots,\mathit{Fdpat}_n\}\rceil$

$\lceil v\ \textbf{isa}\ p\ \ \mathtt{=>}\ \ \{f_1 = v_1\ @\ S_1,\ldots,f_n = v_n\ @\ S_n\}\rceil$

$\quad\Longrightarrow\ \ v\ \textbf{isa}\ p\ \ \mathtt{=>}\ \ \{f_1 = v_1',\ldots,f_n = v_n'\}\ \textbf{and}\bigwedge_{i=1}^{n}\{\lceil v.f_i\ \textbf{isa}\ S_i\rceil\}$

$\lceil v\ \textbf{isa}\ p\ \{f_1 = v_1,\ldots,f_n = v_n\}\rceil$

$\quad\Longrightarrow\ \ p(v)\ \ \mathtt{=>}\ \ \{f_1 = v_1,\ldots,f_n = v_n\}$

Compound predicate abstractions

$\lceil v\ \textbf{isa}\ \textbf{not}\ \mathit{Pname}\{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil$

$\quad\Longrightarrow\ \ \textbf{not}\ \lceil v\ \textbf{isa}\ \mathit{Pname}\{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil$

$\lceil v\ \textbf{isa}\ \mathit{Pname}_1 \mid \mathit{Pname}_2\{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil$

$\quad\Longrightarrow\ \ \lceil v\ \textbf{isa}\ \mathit{Pname}_1\{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil\ \textbf{or}\ \lceil v\ \textbf{isa}\ \mathit{Pname}_2\{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil$

$\lceil v\ \textbf{isa}\ \mathit{Pname}_1\ \&\ \mathit{Pname}_2\{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m,\ldots,\mathit{Fdpat}_n\}\rceil\qquad \mathit{Fdpat}_i$, $1 \leq i \leq m$, is generated by Pname_1

$\quad\Longrightarrow\ \ \lceil v\ \textbf{isa}\ \mathit{Pname}_1\{\mathit{Fdpat}_1,\ldots,\mathit{Fdpat}_m\}\rceil\ \textbf{and}\ \lceil v\ \textbf{isa}\ \mathit{Pname}_2\{\mathit{Fdpat}_{m+1},\ldots,\mathit{Fdpat}_n\}\rceil$

Classifiers

$$\left\lceil \begin{array}{l} \textbf{classify}(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \\ \quad \textbf{as } c_1 \textbf{ when } P_1 \textbf{ return } \{f_{1,1} := w_{1,1}, \ldots, f_{1,m_1} := w_{1,m_1}\} \\ \quad \vdots \\ \quad \textbf{as } c_n \textbf{ when } P_n \textbf{ return } \{f_{n,1} := w_{n,1}, \ldots, f_{n,m_n} := w_{n,m_n}\} \\ \quad \textbf{as } c_{n+1} \textbf{ otherwise return } \{f_{n+1,1} := w_{n+1,1}, \ldots, f_{n+1,m_{n+1}} := w_{n+1m_{n+1}}\} \end{array} \right\rceil$$

$$\Longrightarrow \left\lceil \begin{array}{l} \textbf{predicate } c_1(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \textbf{ when } P_1 \\ \quad \textbf{return } \{f_{1,1} := w_{1,1}, \ldots, f_{1,m_1} := w_{1,m_1}\}; \end{array} \right\rceil$$

$$\lceil \textbf{predicate } d_1(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \textbf{ when } P_1; \rceil$$

$$\left\lceil \begin{array}{l} \textbf{predicate } c_2(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \textbf{ when } P_2 \textbf{ and not } d_1(v_1, \ldots, v_m) \\ \quad \textbf{return } \{f_{2,1} := w_{2,1}, \ldots, f_{2,m_2} := w_{2,m_2}\}; \end{array} \right\rceil$$

$$\lceil \textbf{predicate } d_2(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \textbf{ when } d_1(v_1, \ldots, v_m) \textbf{ or } P_2; \rceil$$

$$\ldots$$

$$\left\lceil \begin{array}{l} \textbf{predicate } c_n(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \textbf{ when } P_n \textbf{ and not } d_{n-1}(v_1, \ldots, v_m) \\ \quad \textbf{return } \{f_{n,1} := w_{n,1}, \ldots, f_{n,m_n} := w_{n,m_n}\}; \end{array} \right\rceil$$

$$\lceil \textbf{predicate } d_n(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \textbf{ when } d_{n-1}(v_1, \ldots, v_m) \textbf{ or } P_n; \rceil$$

$$\left\lceil \begin{array}{l} \textbf{predicate } c_{n+1}(v_1 \mathbin{@} S_1, \ldots, v_m \mathbin{@} S_m) \textbf{ when not } d_n(v_1, \ldots, v_m) \\ \quad \textbf{return } \{f_{n+1,1} := w_{n+1,1}, \ldots, f_{n+1,m_{n+1}} := w_{n+1,m_{n+1}}\}; \end{array} \right\rceil$$

B Bindings escaping "or"

In the static and dynamic semantics presented in Section 3, bindings never escape from **or** predicate expressions. Relaxing this constraint provides extra convenience to the programmer and permits more values to be reused rather than recomputed. It is also equivalent to permitting overloaded predicates or multiple predicate definitions — so far we have permitted only a single definition of each predicate.

For example, the two `ConstantFold` methods of Section 2.3 can be combined into a single method. Eliminating code duplication is a prime goal of object-oriented programming, but the previous version repeated the body twice.

```
-- handle case of adding zero to anything (but don't be ambiguous
-- with existing method for zero plus a constant)
method ConstantFold(e@BinopExpr{ op@IntPlus, arg1=a1, arg2=a2 })
  when (a1@IntConst{ value=v } and test(v==0) and not(a2@IntConst) and let res := a2)
    or (a2@IntConst{ value=v } and test(v==0) and not(a1@IntConst) and let res := a1) {
  ... -- increment counter, or do other common work here
  return res; }
```

As another example, the `LoopExit` example of Section 2.4 can be extended to present a view which indicates which branch of the `CFG_2succ` is the loop exit and which the backward branch. When performing iterative dataflow, this is the only information of interest, and in our current implementation (which uses predicate classes [Cha93b] we generally recompute this information after discovering that an object is a `LoopExit`. Presenting a view which includes this information improves the code's readability and efficiency.

```
predsignature LoopExit(CFGnode)
  return { loop:CFGloop, next_looping:CFGedge, next_exiting:CFGedge };
predicate LoopExit(n@CFG_2succ{ next_true: t, next_false: f })
  when (test(t.is_loop_exit) and let nl := t and let ne := f)
    or (test(f.is_loop_exit) and let nl := f and let ne := t)
  return { loop := nl.containing_loop, next_looping := nl, next_exiting := ne };
```

Permitting bindings which appear on both sides of **or** to escape requires the following changes to the dynamic semantics of Figure 3:

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{true}, K'\rangle}{\langle P \textbf{ or } Q, K\rangle \Rightarrow \langle \mathit{true}, K'\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{false}, K'\rangle \qquad \langle Q, K\rangle \Rightarrow \langle \mathit{true}, K''\rangle}{\langle P \textbf{ or } Q, K\rangle \Rightarrow \langle \mathit{true}, K''\rangle}$$

$$\frac{\langle P, K\rangle \Rightarrow \langle \mathit{false}, K'\rangle \qquad \langle Q, K\rangle \Rightarrow \langle \mathit{false}, K''\rangle}{\langle P \textbf{ or } Q, K\rangle \Rightarrow \langle \mathit{false}, K\rangle}$$

The static semantics of Figure 4 must be modified to add a helper function and to replace a typechecking rule:

$\sqcup_{\text{env}}(\Gamma, \Gamma') = \Gamma''$	Pointwise lub over typing environments. For each $v \in \text{dom}(\Gamma'') = \text{dom}(\Gamma) \cap \text{dom}(\Gamma')$, if $\Gamma \models v : T$ and $\Gamma' \models v : T'$, then $\Gamma'' \models v : T \sqcup T'$.

$$\frac{\Gamma \vdash P_1 \Rightarrow \Gamma' \qquad \Gamma \vdash P_2 \Rightarrow \Gamma'' \qquad \sqcup_{\text{env}}(\Gamma', \Gamma'') = \Gamma'''}{\Gamma \vdash P_1 \textbf{ or } P_2 \Rightarrow \Gamma'''}$$

Finally, canonicalization must account for the new semantics of **or**. In order to permit replacement of variables by their values, we introduce a new compile-time-only ternary conditional operator **?:** for each variable bound on each side of the predicate. The first argument is the predicate expression on the left-hand side of the **or** expression; the second and third arguments are the values on each side of the **or**.

Canonicalizing this new **?:** expression requires ordering the tests canonically; any ordering will do. This may necessitate duplication of some expressions, such as transforming $b?e_1{:}a?e_2{:}e_3$ into $a?(b?e_1{:}e_2){:}(b?e_1{:}e_3)$ so that those two expressions are not considered distinct. With these two modifications, the tautology test is once again sound and complete.