

Ptolemy: A Language of Quantified, Typed Events

Hridesh Rajan and Gary T. Leavens

TR #07-13

Initial Version: July 26, 2007.

Keywords: implicit-invocation languages, aspect-oriented programming languages, quantification, pointcut, join point, context exposure, type checking, event types, event expressions.

CR Categories:

D.3.3 [*Programming Languages*] Language Constructs and Features - Control structures

Copyright © 2007, Hridesh Rajan and Gary T. Leavens.

Department of Computer Science
226 Atanasoff Hall
Iowa State University
Ames, Iowa 50011-1041, USA

Ptolemy: A Language of Quantified, Typed Events

Hridesh Rajan
Iowa State University
hridesh@cs.iastate.edu

Gary T. Leavens
Iowa State University
leavens@cs.iastate.edu

Abstract

This paper defines Ptolemy. The novelty of Ptolemy is the notion of event types and quantification based on event types. We give the syntax, operational semantics and type rules for the language, and discuss its meta-theory.

1. Ptolemy’s Design

In this section, we describe Ptolemy, a language with quantified, typed events that extends implicit invocation (II) languages with ideas from aspect-oriented (AO) languages. Ptolemy features new mechanisms for declaring event types and events. Our description includes syntax, examples, semantics, and type checking rules.

1.1 Overview

Ptolemy is inspired by II languages such as Rapide [8] and AO languages such as AspectJ [5]. It also incorporates some ideas from Eos [11] and Caesar [9]. As a small, core language, its technical presentation shares much in common with Clifton and Leavens’s MiniMAO₁ [1, 2], which itself builds on Classic Java [4] and Featherweight Java [7]. The object-oriented part of Ptolemy follows MiniMAO₀. While it has classes, objects, inheritance, and subtyping, it does not have **super**, interfaces, exception handling, built-in value types, privacy modifiers, or abstract methods. The novel features of Ptolemy are found in its event model and type system. In the syntax these novel features are: an event type declaration (**evtype**), and an event expression (**event**).

Like Eos [11], Ptolemy does not have special syntax for aspects and advice. Instead it has the capability to replace all events in a specified set (a pointcut) with a call to a method. Following II terminology, we call such methods *event handlers* or simply *handlers*. Each handler takes a handler chain as its first argument. The type of the handler chain is the same as the return type of the event type and the return type of all handlers in the chain. A *handler chain* can be evaluated to run any remaining handlers and the original event. Handlers and handler chains produce an effect equivalent to AspectJ’s around advice and **proceed** constructs.

Like II languages a Ptolemy module can register to receive event notifications. This capability is the same as “deployment” in the AO languages Eos and Caesar [9]. However, like II languages, registration makes explicit the receiver instance that will run the handler, and allows instance-level advising features to be easily programmed [12]. Singleton “aspects” that are created at the start

of the program and automatically registered can also be easily programmed or added as a syntactic sugar.

1.2 Syntax

Ptolemy’s syntax is shown in Figure 1 and explained below. A program consists of a sequence of declarations followed by an expression. The expression can be thought of as the body of a (static, or receiverless) “main” method. We next explain declarations, pointcut descriptions, and expressions.

```

prog ::= decl* e
decl ::= class c extends c { field* meth* binding* }
      | evtype p { form* }
field ::= t f;
meth ::= t m (form*) { e }
binding ::= t around (form*) pcd : m
form ::= t var, where var ≠ this
pcd ::= p | cflow (pcd) | pcd && pcd | pcd ‘|’ pcd
e ::= new c () | var | null | e.m (e*) | e.f | e.f = e
    | cast t e | form = e; e
    | register (e) | event p { e }

c, d ∈ C, the set of class names
p ∈ P, the set of evtype names
t, s, u ∈ T, the set of types, where (C ∪ P) ⊆ T
f ∈ F, the set of field names
m ∈ M, the set of method names
var ∈ {this} ∪ V, V is the set of variable names

```

Figure 1. Ptolemy’s Abstract syntax, based on [1, Figure 3.1, 3.7].

1.2.1 Declarations

There are only two declaration forms that may appear at the top level of a Ptolemy program: classes and event type declarations. These may not be nested. A class has exactly one superclass, named in its **extends** clause. It may declare several fields (*field**), methods (*meth**), and bindings (*binding**). Field declarations are written, as in C++ and Java, with a type expression followed by a field name. Methods also have a C++ or Java-like syntax, although their body is an expression. An example is given in Figure 2. As in Eos, bindings associate a set of events, described by a pointcut description (PCD), to a method. An example is shown in Figure 2, which contains a binding on line 23. The binding tells Ptolemy to run the method `update` whenever events of type `FEChange` are executed.

An event type (**evtype**) declaration has a return type (*t*), a name (*p*), and zero or more context variable declarations (*form**). These context declarations specify the types and names of reflective information exposed by conforming events. An example is given in Figure 2 on line 17. In writing examples of event types, as in Figure 2, we show each formal parameter declaration (*form*) as terminated by a semicolon (;). In examples showing the declarations of methods and bindings, we use commas to separate each *form*. The intention of this event type declaration is to provide a named abstraction for a set of events, with result type `FEElement`, that con-

```

1 class FElement extends Object{}
2 class Point extends FElement { /* ... */
3   Number x; Number y;
4   FElement setX(Number x) {
5     Point changedFE = this;
6     event FEChange {
7       this.x = x; this
8     }
9   }
10  FElement makeEqual(Point other) {
11    Point changedFE = other;
12    event FEChange {
13      other.x = this.x; other.y = this.y; other
14    }
15  }
16 }
17 FElement evtype FEChange { FElement changedFE; }
18 class Update extends Object {
19   FElement last;
20   Update init() { register(this) }
21   FElement update(FElement next, FElement changedFE)
22   { this.last = changedFE; next }
23   FElement around(FElement changedFE) FEChange:update
24 }

```

Figure 2. Drawing Editor in Ptolemy

tribute to an abstract state change in a figure element, such as moving a point, line, etc. This event type declares only one context variable, `changedFE`, which denotes the `FElement` instance that is being changed. An event can only be of this type if: (a) it has the stated result type and (b) it binds the context variable `changedFE` to some value in the lexical scope surrounding the event shadow, as shown in Figure 2.

1.2.2 Quantification: Pointcut Descriptions

The syntax for pointcut descriptions (or PCDs, sometimes called pointcut designators) has one basic form and three recursive forms, corresponding to basic and complex events in II languages. The basic PCD is the named PCD, which denotes the set of events that are identified by the programmer using `event` expressions with that name. The context exposed by such a named event is the context available at the event identified by the programmer. An example appears in line 23 of Figure 2, which denotes events identified with the type `FEChange`.

The `cflow`, or control flow, PCD is similar to AspectJ’s `cflow` PCD. It names all programmer-identified events that occur during the execution of the PCD it contains, including those named by that PCD. The context exposed by such a `cflow` PCD is the context exposed by the underlying PCD. For example `cflow` (`FEChange`) includes all events in `FEChange`, as well as all those that occur during their execution, and it exposes all the context that `FEChange` exposes. However, unlike AspectJ, in Ptolemy only explicitly identified events that occur in the control flow of `FEChange` are considered to be events, not all possible events of AspectJ’s predefined event kinds. This change makes clear where advice can run.

As in AspectJ, disjunction (`|`) of two PCDs gives the union of the sets of events denoted by the two PCDs. The context exposed by the disjunction is the intersection of the context exposed by the two PCDs. However, if an identifier I is bound in both contexts, then I ’s value in the exposed context is I ’s value from the right hand PCD’s context.

Similarly, the conjunction of two PCDs intersects the set of events denoted by the two PCDs. A conjunction event exposes context that is the union of the context exposed by the two PCDs.

Again, if an identifier I is bound in both contexts, then I ’s value in the exposed context is I ’s value from the right hand PCD’s context.

1.2.3 Expressions

Ptolemy is an expression language. Thus the syntax for expressions includes several standard object-oriented (OO) expressions and also some expressions that are specific to aspects.

The standard OO expressions include object construction (`new c()`), variable dereference (`var`, including `this`), field dereference (`e.f`), `null`, cast (`cast t e`), imperative such as assignment to a field (`e1.f = e2`), and a definition block (`t var = e1; e2`). Their semantics and typing is fairly standard [1, 2]. For convenience in examples, we use sequencing, written $e_1; e_2$, as a syntactic sugar for a definition block, which desugars as follows: $e_1; e_2 \Rightarrow t_1 var_1 = e_1; e_2$, where t_1 is the static type of e_1 and var_1 is fresh.

There are also two expressions that are new in Ptolemy: `register` and `event`. The expression `register(e)` evaluates e to an object o , puts o into the list of active objects, and returns o . For example line 20 of Figure 2 is a method that, when called, will make the method’s receiver (`this`) active. The expression `event p {e}` declares the expression e as an event of type p and runs any handler methods of active objects that is applicable to p . That is, it marks e as the shadow [6] of an event of type p . The type named, p , must be an event type. This type name: (i) identifies the event for purposes of quantification, much like an annotation would in AspectJ 5, and (ii) is used in type checking.

Handler methods are called with an active object as a receiver. The call passes a handler chain as its first argument. This handler chain can be evaluated to run any remaining handlers, or to run the body of the event expression. Since Ptolemy has handler chains bound to formal parameters, there is no need for a special `proceed` syntax. Instead, we often use `proceed` as an identifier so that mentioning it in a Ptolemy program has a similar effect to `proceed` in AspectJ.

An example demonstrating these features is shown in Figure 2. The event declared on lines 6–8 has a body consisting of the sequence expression on line 7. Notice that the body of the `setX` method contains a block expression, where the definition on line 5 binds `this` to `changedFE`, and then evaluates its body, the event expression. This definition makes the value of `this` available in the context variable `changedFE`, which is needed by the event type `FEChange`. In this figure, the event declared on lines 12–14 encloses the sequence expression on line 13. As required by the event type, the definition on line 11 of Figure 2 makes the value of `other` available in the context variable `changedFE`. Thus the first and the second event expressions are given different bindings for the context variable `changedFE`, however, code that advises this event type will be able to access this context variable uniformly using the name `changedFE`.

The evaluation of an event expression first looks for any applicable bindings for objects in the active (registered) list. The handler methods from such applicable bindings are formed into a list, ordered in reverse of the order of object activation, with the most recently registered object’s handlers first. This semantics is designed to allow more recently registered objects to control whether previously registered objects have their advice run, by evaluating a handler chain (or not). Similarly, within an object’s handlers, subclass and textually later bindings are allowed the same control over superclass and textually earlier bindings. That is, when handler methods from applicable bindings for an object are formed into a list, handlers from that subclasses of that object’s type appear before handlers declared in its superclasses. Furthermore, for bindings declared in the same class, handlers for textually earlier bindings appear after handlers for later bindings.

Consider a Ptolemy program that combines Figure 2 followed by the main expression

```
Update u = new Update().init();
Point p = new Point();
p.setX(new Zero());
u.last
```

This main expression creates and registers an `Update` object, which it names `u`. It then creates a `Point` object, and binds it to `p`. The call to `setX` binds the formal `x` to the object representing the number 0, and then runs the body of `setX`. Since the body contains an event expression, and since there is an active object (`u`) that contains a binding for that event, that binding’s handler method is run. This method, `update`, is called with receiver `u`, a handler chain as the first argument, and the value of `changedFE` as the second argument. The body of `update` assigns to `u`’s field `last`, and then calls the handler chain (the expression `next`), which executes the body of the event expression (starting at line 9 of Figure 2) in its original environment. The body of the event expression returns the value of `this`, which, since the environment of the call to `setX` has been restored, is the value of `p`. This value is returned as the value of the handler chain, and hence as the result of the method `update`. In turn, this result, `p`, is used as the value of the event expression, and hence as the value of the call to `setX`. Thus the expression in the last line of the main program’s expression, `u.last` denotes the same object as `p`.

The grammar only allows one event type to be named in an event expression. However, it is convenient to allow a list of event types as a syntactic sugar. The desugaring to a nest of event expression is as follows.

$$\begin{aligned} & \mathbf{event} \ p_1, \dots, p_n \ \{e\} \\ \Rightarrow & \mathbf{event} \ p_1 \{ \dots \mathbf{event} \ p_n \ \{e\} \dots \} \end{aligned}$$

Note, however, that this sugar does not make the events listed occur simultaneously; they instead occur in a definite order.

1.3 Operational Semantics

This section defines a small step operational semantics for Ptolemy. This semantics has been implemented in the logic programming language λ Prolog, using the Teyjus system [10]. This semantics and its description in this section is adapted from Clifton’s work [1, 2, 3], which builds on Classic Java [4]. Following these works, a program’s declarations are simply formed into a fixed list, which is used in the semantics of expressions. The small steps of the operational semantics thus gives a semantics of programs by giving a semantics of expressions.

The expression semantics relies on five expressions that are not part of Ptolemy’s surface syntax. These expressions allow the semantics to record final or intermediate states of the computation, and are shown in Figure 3. The `loc` expression represents locations in the store. The `under` expression is used as a way to mark when the evaluation stack needs popping. The `chain` expression is used to represent handler chains. It contains an ordered list of handler records (H), a PCD type (θ), an expression (e), an environment (ρ), and a type environment (Π). The type θ and the type environment Π (see Figure 9) are not used by the operational semantics, but only in the type soundness proof. Each handler record (h) contains the information necessary to call a handler method: a value that will be the receiver object of the method call (v), a method name (m), and an environment (ρ). The environment maps identifiers to values (v), and is used to assemble the method call arguments. The environment recorded at the top level of the chain expression is used to run the expression e in the chain expression. The two exceptions record various problems orthogonal to the type system.

The small steps taken in the semantics transition from one configuration to another. These configurations are described in Fig-

$$\begin{aligned} e & ::= \mathit{loc} \mid \mathbf{under} \ e \mid \mathbf{chain} \ (H, \theta) \ (e, \rho, \Pi) \\ & \quad \mid \mathit{NullPointerException} \mid \mathit{ClassCastException} \\ H & ::= h + H \mid \bullet \\ h & ::= \langle \mathit{loc}, m, \rho \rangle \\ \rho & ::= \{j : v_k\}_{k \in K}, \quad \text{where } K \text{ is finite, } K \subseteq I \\ v & ::= \mathit{loc} \mid \mathbf{null} \end{aligned}$$

$\mathit{loc} \in \mathcal{L}$, the set of locations

Figure 3. Added syntax for Ptolemy’s operational semantics.

ure 4. A configuration contains an expression (e), a stack (J), a store (S), and an ordered list of active objects (A). Stacks are an ordered list of frames, each frame recording the static environment (ρ) and some other information. (The type environments Π are only used in the type soundness proof.) There are four types of stack frame. Method execution frames (**execframe**) record the name of the method invoked.¹ Event frames (**evframe**) record the name of the event type being run. Definition frames (**defframe**) are used for extending the environment with an identifier to value mapping, and continuation frames (**contframe**) are used when an event’s body is being run. Stores are maps from locations to objects. Objects map field names to values. A value is a location or **null**.

Domains:

$\Gamma ::= \langle e, J, S, A \rangle$	“Configurations”
$J ::= \nu + J \mid \bullet$	“Stacks”
$\nu ::=$	“Frames”
$\mathbf{execframe} \ m \ \rho \ \Pi$	“Method execution”
$\mathbf{defframe} \ \rho \ \Pi$	“Definition”
$\mathbf{evframe} \ p \ \rho \ \Pi$	“Event execution”
$\mathbf{contframe} \ \rho \ \Pi$	“Continuation”
$S ::= \{ \mathit{loc}_k \mapsto o_k \}_{k \in K},$	“Stores”
where K is finite	
$o ::= [c, F]$	“Object Records”
$F ::= \{f_k \mapsto v_k\}_{k \in K},$	“Field Maps”
where K is finite	
$A ::= \mathit{loc} + A \mid \bullet$	“Active List”

Figure 4. Domains used in the semantics, based on [1].

As is usual [14] the semantics is presented as a set of evaluation contexts \mathbb{E} and a one-step reduction relation that acts on the position in the overall expression identified by the evaluation context. This two-part presentation avoids the need for writing out standard recursive rules and has the advantage of more clearly presenting the order of evaluation. Figure 5 defines evaluation contexts, and hence the order of evaluation for Ptolemy. The language uses a strict leftmost, innermost evaluation policy, which thus uses call-by-value. The initial configuration for a program with main expression e is $\langle \mathbf{under} \ e, (\mathbf{execframe} \ \mathit{main} \ \{\} \ \{\}) + \bullet, \{\}, \bullet \rangle$, which starts evaluation of e in a frame with an empty environment, and with an empty store and list of active objects.

Evaluation contexts:

$$\begin{aligned} \mathbb{E} & ::= - \mid \mathbb{E}.m(e\dots) \mid v.m(v\dots\mathbb{E}\dots) \mid \mathbf{cast} \ t \ \mathbb{E} \\ & \quad \mid \mathbb{E}.f \mid \mathbb{E}; e \mid \mathbb{E}.f=e \mid v.f=\mathbb{E} \mid t \ \mathit{var}=\mathbb{E}; e \\ & \quad \mid \mathbf{register} \ (\mathbb{E}) \mid \mathbf{under} \ \mathbb{E} \end{aligned}$$

Figure 5. Evaluation contexts for Ptolemy, based on [1].

Figure 6 presents the operational semantics of Ptolemy. In these rules all of the hypotheses are really side conditions and side definitions for use in the rule. The rules all make implicit use

¹ The name of the method being run is not needed in the present language, but could be used in AspectJ-style call PCDs.

$$\begin{array}{c}
\text{Evaluation relation: } \hookrightarrow: \Gamma \rightarrow \Gamma \\
\text{(NEW)} \\
\frac{loc \notin \text{dom}(S)}{S' = S \oplus (loc \mapsto [c. \{f \mapsto \mathbf{null} \mid f \in \text{dom}(\text{fieldsOf}(c))\}])} \\
\langle \mathbb{E}[\mathbf{new } c()], J, S, A \rangle \hookrightarrow \langle loc, J, S', A \rangle \\
\\
\text{(VAR)} \\
\frac{\rho = \text{envOf}(\nu) \quad v = \rho(\text{var})}{\langle \mathbb{E}[\text{var}], \nu + J, S, A \rangle \hookrightarrow \langle \mathbb{E}[v], \nu + J, S, A \rangle} \\
\text{(GET)} \\
\frac{[c. F] = S(loc) \quad v = F(f)}{\langle \mathbb{E}[loc.f], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[v], J, S, A \rangle} \\
\\
\text{(SET)} \\
\frac{[c. F] = S(loc) \quad S' = S \oplus (loc \mapsto [c. F \oplus (f \mapsto v)])}{\langle \mathbb{E}[loc.f = v], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[v], J, S', A \rangle} \\
\\
\text{(DEF)} \\
\frac{\rho = \text{envOf}(\nu) \quad \rho' = \rho \oplus (\text{var} \mapsto v) \quad \Pi = \text{tenvOf}(\nu) \quad \Pi' = \Pi \uplus \{\text{var} : t\} \quad \nu' = \mathbf{defframe } \rho' \Pi'}{\langle \mathbb{E}[t \text{ var} = v; e], \nu + J, S, A \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } e], \nu' + \nu + J, S, A \rangle} \\
\\
\text{(CALL)} \\
\frac{\begin{array}{l} [c. F] = S(loc) \quad \nu = \mathbf{execframe } m \rho \Pi \\ t \ m(t_1 \text{ var}_1, \dots, t_n \text{ var}_n) \{e\} = \text{methodBody}(c, m) \\ \rho = \{\text{var}_i \mapsto v_i \mid 1 \leq i \leq n\} \oplus (\mathbf{this} \mapsto loc) \\ \Pi = \{\text{var}_i : t_i \mid 1 \leq i \leq n\} \uplus \{\mathbf{this} : c\} \end{array}}{\langle \mathbb{E}[loc.m(v_1, \dots, v_n)], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } e], \nu + J, S, A \rangle} \\
\\
\text{(CAST)} \\
\frac{[c. F] = S(loc) \quad c \preccurlyeq t}{\langle \mathbb{E}[\mathbf{cast } t \ \text{loc}], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[loc], J, S, A \rangle} \\
\text{(NCAST)} \\
\frac{}{\langle \mathbb{E}[\mathbf{cast } t \ \mathbf{null}], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{null}], J, S, A \rangle} \\
\\
\text{(SKIP)} \\
\frac{}{\langle \mathbb{E}[v; e], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[e], J, S, A \rangle} \\
\text{(UNDER)} \\
\frac{}{\langle \mathbb{E}[\mathbf{under } v], \nu + J, S, A \rangle \hookrightarrow \langle \mathbb{E}[v], J, S, A \rangle} \\
\\
\text{(REGISTER)} \\
\frac{}{\langle \mathbb{E}[\mathbf{register } loc], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[loc], J, S, loc + A \rangle} \\
\\
\text{(EVENT)} \\
\frac{\begin{array}{l} \rho = \text{envOf}(\nu) \quad (t \ \mathbf{evtype } p \{t_1 \ \text{var}_1, \dots, t_n \ \text{var}_n\}) \in CT \\ \rho' = \{\text{var}_i \mapsto v_i \mid \rho(\text{var}_i) = v_i\} \quad \Pi = \text{tenvOf}(\nu) \\ \nu' = \mathbf{evframe } \rho' \Pi \quad H = \text{hbind}(\nu' + \nu + J, S, A) \\ \pi = \{\text{var}_i : t_i \mid 1 \leq i \leq n\} \quad \theta = \mathbf{pcd } t, \pi \end{array}}{\langle \mathbb{E}[\mathbf{event } p \{e\}], \nu + J, S, A \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } (\mathbf{chain } (H, \theta) (e, \rho, \Pi))], \nu' + \nu + J, S, A \rangle} \\
\\
\text{(CHAIN-DONE)} \\
\frac{\nu = \mathbf{contframe } \rho \Pi}{\langle \mathbb{E}[\mathbf{chain } (\bullet, \theta) (e, \rho, \Pi)], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } e], \nu + J, S, A \rangle} \\
\\
\text{(CHAIN-RUN)} \\
\frac{\begin{array}{l} [c. F] = S(loc) \quad \nu = \mathbf{execframe } m \rho''' \Pi' \quad \theta = \mathbf{pcd } t_1, \pi \\ t_1 \ m(t_1 \ \text{var}_1, \dots, t_n \ \text{var}_n) \{e'\} = \text{methodBody}(c, m) \\ n \geq 1 \quad \rho'' = \{\text{var}_i \mapsto v_i \mid 2 \leq i \leq n, v_i = \rho(\text{var}_i)\} \\ \rho''' = \rho'' \oplus \{\text{var}_1 \mapsto (\mathbf{chain } (H, \theta) (e, \rho', \Pi))\} \oplus \{\mathbf{this} \mapsto loc\} \\ \Pi' = \{\text{var}_i : t_i \mid 1 \leq i \leq n\} \uplus \{\mathbf{this} : c\} \end{array}}{\langle \mathbb{E}[\mathbf{chain } ((\langle loc, m, \rho \rangle + H), \theta) (e, \rho', \Pi)], J, S, A \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } e'], \nu + J, S, A \rangle}
\end{array}$$

Figure 6. Operational semantics of Ptolemy, based on [1].

of a fixed (global) list of the program’s declarations. The rules use \oplus as an overriding operator for finite functions. That is, if $S' = S \oplus (loc \mapsto v)$, then $S'(loc') = v$ if $loc' = loc$ and otherwise $S'(loc') = S(loc')$. The *fieldsOf* function uses the class table to determine the list of field declarations for a given class

(and its superclasses), considered as a mapping from field names to their types. Thus, (NEW) says that the store is updated to map a fresh location to an object of the given class that has each of its fields set to null.

In the (VAR) rule, *envOf*(ν) returns the environment from the current frame ν , ignoring any other information in ν .

$$\begin{array}{l}
\text{envOf}(\mathbf{execframe } m \rho \Pi) = \rho \\
\text{envOf}(\mathbf{defframe } \rho \Pi) = \rho \\
\text{envOf}(\mathbf{evframe } p \rho \Pi) = \rho \\
\text{envOf}(\mathbf{contframe } m \rho \Pi) = \rho
\end{array}$$

Thus the (VAR) rule says that the value of a variable, including **this**, is simply looked up in the environment of the current frame. The (CALL) rule implements dynamic dispatch by looking up the method m starting from the dynamic class (c) of the receiver object (loc), looking in superclasses if necessary, using the auxiliary function *methodBody* (not shown here). The body is executed in an **execframe** with an environment that binds the methods formals, including **this**, to the actual parameters. Since methods do not nest, and since expressions access object fields by starting from an explicit object there is no other context available to a method. Note that **under** e is used in the resulting configuration for the (CALL) rule. This expression is used whenever a new frame is pushed on the stack, to record that the stack should be popped when the evaluation of e is finished. The (UNDER) rule pops the stack when evaluation of its subexpression is finished. The (GET) and (SET) rules are standard. The value of a field assignment is the value being assigned. The (CAST) rule simply checks that the dynamic class of the object is a subtype of the type given in the expression. The (NCAST) rule allows **null** to be cast to any type. The (DEF) rule allows for local definitions. It is similar to **let** in other languages, but with a more C++ and Java-like syntax. It simply binds the variable given to the value in an extended environment. Since a new frame is pushed on the stack, the body, e , is evaluated inside an “under” expression, which pops the stack when e is finished.

The (REGISTER) rule simply puts the object being activated at the front of the list of active objects. The bindings in this object are thus given control before others already in the list. Notice that an object can appear in this list multiple times.

The (EVENT) rule is central to Ptolemy’s semantics, as it starts the running of handler methods. In essence, the rule forms a new frame for running the event, and then looks up bindings applicable to the new stack, store, and list of active objects. The resulting list of handler records (H) is put into a handler chain (**chain** ($H \theta$)($e \rho \Pi$)), which will execute the handler methods, if any, before the body of the event expression (e) is evaluated. Since a new (event) frame is pushed on the stack, this handler chain is placed inside an **under** expression, so that the stack will be popped when the handler chain is finished.

The auxiliary function *hbind*, defined in Figure 7 uses the program’s declarations, the stack, store, and the list of active objects to produce a list of handler records that are applicable for the event in the current state. When called by the (EVENT) rule, the stack passed to it has a new frame on top that represents the current event.

The *hmatch* function determines, for a particular object loc , what bindings declared in the class of the object referred to by loc are applicable. It looks up the location loc in the store, extracts the class of the object loc refers to, and uses that class to obtain a list of potential bindings. This list is filtered using *match*, which relies on *mpcd* to match a PCD against a particular event on the stack. Each matching binding generates a handler record, recording the active object (which will act as a receiver when the handler method is called), the handler method’s name, and an environment. The environment is obtained by *mpcd*, ultimately from the environment in frames of type **evframe**. This environment is also restricted

$$\begin{aligned}
& hbind(J, S, \bullet) = \bullet \\
& hbind(J, S, loc + A) \\
& = concat(hmatch(CT, J, S, loc), hbind(J, S, A)) \\
& \text{where } CT \text{ is the program's list of declarations} \\
& \text{and } concat(\bullet, H') = H' \\
& \quad concat(h + H, H') = h + concat(H, H') \\
\\
& hmatch(CT, J, S, loc) = match(H, J, S, loc) \\
& \text{where } S(loc) = [c.F] \text{ and } bindings(CT, c) = H \\
\\
& bindings(CT, c) = binds(CT, CT, c) \\
& binds(CT, \bullet, c) = \bullet \\
& binds(CT, ((t \mathbf{evtype} p\{\dots\}) + CT'), c) = binds(CT, CT', c) \\
& binds(CT, ((\mathbf{class} c \mathbf{extends} c' \dots binding_1 \dots binding_n) + CT'), c) \\
& = concat((binding_n + \dots + binding_1 + \bullet), binds(CT, CT', c)) \\
\\
& match(\bullet, J, S, loc) = \bullet \\
& match(binding + H, J, S, loc) \\
& = \mathbf{if} \ mpcd(pcd, J, S) \neq \perp \\
& \quad \mathbf{then} \ \mathbf{let} \ \rho = mpcd(pcd, J, S) \\
& \quad \quad \mathbf{in} \ \mathbf{let} \ \rho' = \{var_i \mapsto \rho(var_i) \mid 1 \leq i \leq n\} \\
& \quad \quad \mathbf{in} \ ((loc, m, \rho') + match(H, J, S, loc)) \\
& \quad \mathbf{else} \ match(H, J, S, loc) \\
& \text{where } binding = t_1 \mathbf{around}(t_1 \ var_1, \dots, t_n \ var_n) pcd : m \\
\\
& mpcd(p, (\mathbf{evframe} p' \rho) + J, S) = \mathbf{if} \ p \equiv p' \ \mathbf{then} \ \rho \ \mathbf{else} \ \perp \\
& mpcd(\mathbf{cflow} pcd, \bullet, S) = \perp \\
& mpcd(\mathbf{cflow} pcd, \nu + J, S) \\
& = \mathbf{if} \ mpcd(pcd, \nu + J, S) \neq \perp \ \mathbf{then} \ mpcd(pcd, \nu + J, S) \\
& \quad \mathbf{else} \ mpcd(pcd, J, S) \\
& mpcd(pcd_1 \ \&\& \ pcd_2, J, S) \\
& = \mathbf{if} \ mpcd(pcd_1, J, S) \neq \perp \wedge mpcd(pcd_2, J, S) \neq \perp \\
& \quad \mathbf{then} \ mpcd(pcd_1, J, S) \uplus mpcd(pcd_2, J, S) \\
& \quad \mathbf{else} \ \perp \\
& mpcd(pcd_1 \ || \ pcd_2, J, S) \\
& = \mathbf{if} \ mpcd(pcd_1, J, S) \neq \perp \wedge mpcd(pcd_2, J, S) \neq \perp \\
& \quad \mathbf{then} \ mpcd(pcd_1, J, S) \uplus mpcd(pcd_2, J, S) \\
& \quad \mathbf{else} \ \mathbf{if} \ mpcd(pcd_1, J, S) \neq \perp \ \mathbf{then} \ mpcd(pcd_1, J, S) \\
& \quad \mathbf{else} \ mpcd(pcd_2, J, S) \\
\\
& \rho \uplus \rho' = \{var \mapsto \rho(var) \mid var \in dom(\rho) \wedge var \notin dom(\rho')\} \cup \rho' \\
& \rho \uplus \rho' = \{var \mapsto \rho'(var) \mid var \in dom(\rho) \wedge var \in dom(\rho')\}
\end{aligned}$$

Figure 7. Auxiliary functions for matching bindings.

to contain just those mappings that are for names in the declared formals of the binding.

When a PCD matches the given stack and store, $mpcd$ returns an environment, otherwise it returns \perp . For named events that match, it returns the environment from the top frame on the stack. For a \mathbf{cflow} PCD, it searches the stack and returns the first environment that matches the enclosed PCD. The disjunction and conjunction PCDs produce an environment that favors their right argument's mappings. For disjunction the result is a kind of intersection, and for conjunction the result is a kind of union.

The evaluation of handler chains is done by the two chain rules. The (CHAIN-DONE) rule handles the case where there are no (more) handler records. It simply runs the event's body expression (e) in the environment (ρ) that was being remembered for it by the handler chain. The environment is made active by using a \mathbf{conf} frame containing it as the top frame on the stack. The expression is put inside an \mathbf{under} expression, so that this new frame will be popped when its evaluation is over.

The (CHAIN-RUN) rule handles the case where there are handler records still to be run. It makes a call to the active object (referred to by loc) in the first handler record, using the method name and environment stored in that handler record. The active object is the receiver of the method call. The first formal parameter is bound to

$$\begin{aligned}
& \text{(NCALL)} \\
& \langle \mathbb{E}[\mathbf{null}.m(v_1, \dots, v_n)], J, S, A \rangle \\
& \hookrightarrow \langle \text{NullPointerException}, \bullet, S, A \rangle \\
\\
& \text{(NGET)} \\
& \langle \mathbb{E}[\mathbf{null}.f], J, S, A \rangle \\
& \hookrightarrow \langle \text{NullPointerException}, \bullet, S, A \rangle \\
\\
& \text{(NSET)} \\
& \langle \mathbb{E}[\mathbf{null}.f = v], J, S, A \rangle \\
& \hookrightarrow \langle \text{NullPointerException}, \bullet, S, A \rangle \\
\\
& \text{(XCAST)} \\
& \frac{[c.F] = S(loc) \quad c \not\approx t}{\langle \mathbb{E}[\mathbf{cast} \ t \ loc], J, S, A \rangle} \\
& \hookrightarrow \langle \text{ClassCastException}, \bullet, S, A \rangle \\
\\
& \text{(NREGISTER)} \\
& \langle \mathbb{E}[\mathbf{register} \ \mathbf{null}], J, S, A \rangle \\
& \hookrightarrow \langle \text{NullPointerException}, \bullet, S, A \rangle
\end{aligned}$$

Figure 8. Operational semantics of expressions that produce exceptions, based on [1].

a handler chain that would run the rest of the handler records (and the original event's body).

The operational semantics rules that result in exceptions are given in Figure 8. These treat some uses of null values and bad casts as exceptions, following Java. Encountering one of these exceptions does not make the semantics be “stuck” and hence the situations that lead to these exceptions are not considered to be type errors. However, all of the resulting configurations are terminal.

1.4 Type Checking

Type checking uses the type attributes defined in Figure 9. (These use some of the notation and ideas from Schmidt's book [13].)

$$\begin{aligned}
\theta ::= & \text{“type attributes”} \\
& \text{OK} \quad \text{“program/top-level decl.”} \\
& \text{OK in } c \quad \text{“method, binding”} \\
& | \ \mathbf{var} \ t \quad \text{“var/formal/field”} \\
& | \ \mathbf{exp} \ t \quad \text{“expression”} \\
& | \ \mathbf{pcd} \ t, \pi \quad \text{“pcd/handler chain”} \\
\pi, \Pi ::= & \{I : \theta_I\}_{I \in K}, \quad \text{“type environments”} \\
& \text{where } K \text{ is finite, } K \subseteq (\mathcal{L} \cup \mathcal{F} \cup \mathcal{M} \cup \mathcal{T} \cup \{\mathbf{this}\}) \cup \mathcal{V}
\end{aligned}$$

Figure 9. Type attributes.

The type checking rule themselves are shown in Figure 10 and 11. See Clifton's thesis [1] for details on these straightforward rules for standard OO expressions. Some rules we use the overriding union notation \uplus , defined in Figure 7 [13]. As in Clifton's work [1, 2], the type checking rules are stated using a fixed class table (list of declarations) CT , which can be thought of as an implicit (hidden) inherited attribute. This class table is used implicitly by many of the auxiliary functions. For ease of presentation, we also follow Clifton in assuming that the names declared at the top level of a program are distinct and that the extends relation on classes is acyclic.

The type checking of PCDs involves their return type and the typing context (a map from variable names to types) that they make available [1]. The return type and typing context of a named PCD are declared where the event type named is declared. For example, the $\mathbf{FEChange}$ PCD has \mathbf{FE} Element as its return type and the typing context that associates $\mathbf{changedFE}$ to the type \mathbf{FE} Element.

$\frac{\text{(NEW EXP TYPE)} \quad \text{isClass}(c)}{\Pi \vdash \mathbf{new} \ c() : \mathbf{exp} \ c}$	$\frac{\text{(CAST EXP TYPE)} \quad \text{isType}(t)}{\Pi \vdash \mathbf{cast} \ t \ e : \mathbf{exp} \ t}$
$\frac{\text{(NULL EXP TYPE)} \quad \text{isClass}(c)}{\Pi \vdash \mathbf{null} : \mathbf{exp} \ c}$	$\frac{\text{(LOC EXP TYPE)} \quad (loc : \mathbf{var} \ t) \in \Pi}{\Pi \vdash loc : \mathbf{exp} \ t}$
$\frac{\text{(GET EXP TYPE)} \quad \Pi \vdash e : \mathbf{exp} \ c \quad \text{fieldsOf}(c)(f) = t}{\Pi \vdash e.f : \mathbf{exp} \ t}$	
$\frac{\text{(SET EXP TYPE)} \quad \Pi \vdash e : \mathbf{exp} \ c \quad \text{fieldsOf}(c)(f) = t \quad \Pi \vdash e' : \mathbf{exp} \ t' \quad t' \preceq t}{\Pi \vdash e.f = e' : \mathbf{exp} \ t'}$	
$\frac{\text{(DEF EXP TYPE)} \quad \text{isType}(t) \quad \Pi \vdash e_1 : \mathbf{exp} \ t_1 \quad t_1 \preceq t \quad \Pi' = \Pi \cup \{\mathbf{var} \ t\} \quad \Pi' \vdash e_2 : t_2}{\Pi \vdash t \ \mathbf{var} = e_1; e_2 : \mathbf{exp} \ t_2}$	
$\frac{\text{(NP EXCEPTION EXP TYPE)}}{\Pi \vdash \mathbf{NullPointerException} : \mathbf{exp} \ \perp}$	
$\frac{\text{(CC EXCEPTION EXP TYPE)}}{\Pi \vdash \mathbf{ClassCastException} : \mathbf{exp} \ \perp}$	

Figure 10. Type-checking rules for OO features.

Since control flow PCDs are dynamic, their return type cannot be used, so we assign them a return type of \top , which is considered a supertype of `Object`, but is not legal as a return type itself. The typing context of a cflow PCD is the typing context of the underlying PCD. Thus if a cflow PCD is not conjoined with any other PCD, the PCD will lead to a type error.

For a disjunction PCD, the return type is the least upper bound of the two PCD's return types, and the typing context is the intersection of the two typing contexts. For each common names I that is in the domain of both contexts, the type exposed for I is the least upper bound of the two types assigned to I by the two PCDs. This makes sense because only one of the two event types may apply.

For the conjunction PCD, the return type is the greatest lower bound of the two PCD's return types, and the typing context is a right-biased overriding union of the two typing contexts. In such a union, each common name I mapped to the type assigned to I by the PCD on right hand side of the conjunction. Note that since a particular PCD must be ultimately based on named event types, and since Ptolemy does not have subtype relationships among named event types, it is usually only sensible to use conjunctions in which one side is not a `cflow` PCD. When this is done, the return type will be that of the named PCD, since the return type of the cflow PCD is \top .

In an `event` expression, the result type of the body expression, e , must be the one declared by the event type, p , and the lexical scope available (at e) must provide the context demanded by p .

In the type checking rules above we use several auxiliary functions. Most of these are taken from Clifton's dissertation [1, Figure 3.3]. Following Clifton, the class table (i.e., the program's list of top-level class and event type declarations) CT is treated as a finite function from type names to their declarations. Thus, for example, in a program that includes Figure 2, $CT(\text{Point})$ would return the declaration of `Point` given in that figure. We exploit this treatment in Figure 12.

$\frac{\text{(CHECK PROGRAM)} \quad (\forall i \in \{1..n\} :: \vdash \text{decl}_i : \text{OK}) \quad \Pi' \vdash e : \mathbf{exp} \ t}{\vdash \text{decl}_1 \dots \text{decl}_n \ e : \mathbf{prog} \ t}$	
$\frac{\text{(CHECK CLASS)} \quad \text{isClass}(c') \quad (\forall j \in \{1..m\} :: \vdash \text{meth}_j \text{OK in } c) \quad (\forall k \in \{1..o\} :: \vdash \text{binding}_k \text{OK in } c) \quad (\forall i \in \{1..n\} :: \text{isClass}(t_i) \wedge f_i \notin \text{dom}(\text{fieldsOf}(c')))}{\vdash \mathbf{class} \ c \ \mathbf{extends} \ c' \ \{t_1 \ f_1; \dots \ t_n \ f_n; \text{meth}_1 \dots \text{meth}_m \ \text{binding}_1 \dots \text{binding}_o\} : \text{OK}}$	
$\frac{\text{(CHECK EVTYPE)} \quad \text{isType}(t)}{\vdash t \ \mathbf{evtype} \ p \ \{t_1 \ \mathbf{var}_1; \dots \ t_n \ \mathbf{var}_n\} : \text{OK}}$	
$\frac{\text{(CHECK METHOD)} \quad \text{isClass}(t) \quad \{var_1 : t_1, \dots, var_n : t_n, \mathbf{this} : c\} \vdash e : \mathbf{exp} \ t' \quad t' \preceq t \quad CT(c) = \mathbf{class} \ c \ \mathbf{extends} \ d \ \{\dots\} \quad \text{override}(m, d, t_1 \times \dots \times t_n \rightarrow t)}{\Pi \vdash t \ m(t_1 \ \mathbf{var}_1, \dots, t_n \ \mathbf{var}_n) \{e\} : \text{OK in } c}$	
$\frac{\text{(CHECK BINDING)} \quad \text{isType}(t) \quad \vdash pcd : \mathbf{pcd} \ t, \pi \quad t \neq \top \quad \text{methodBody}(c, m) = t \ m(t_1 \ \mathbf{var}_1, \dots, t_n \ \mathbf{var}_n) \{e\} \quad n \geq 1 \quad \text{isEVType}(t_1) \quad \{var_2 : t_2, \dots, var_n : t_n\} \subseteq \pi}{\Pi \vdash (t \ \mathbf{around} \ (t_1 \ \mathbf{var}_1, \dots, t_n \ \mathbf{var}_n) \ pcd : m) : \text{OK in } c}$	
$\frac{\text{(EV ID PCD TYPE)} \quad \text{isEVType}(p)}{\vdash p : \mathbf{pcd} \ t, \pi}$	$\frac{\text{(CFLOW PCD TYPE)} \quad \vdash pcd : \mathbf{pcd} \ t, \pi}{\vdash \mathbf{cflow} \ (pcd) : \mathbf{pcd} \ \top, \pi}$
$\frac{\text{(CONJUNCTION PCD TYPE)} \quad \vdash pcd : \mathbf{pcd} \ t, \pi \quad \vdash pcd' : \mathbf{pcd} \ t', \pi' \quad t'' = t \sqcup t' \quad \pi'' = \pi \cup \pi'}{\vdash pcd \ \&\& \ pcd' : \mathbf{pcd} \ t'', \pi''}$	$\frac{\text{(DISJUNCTION PCD TYPE)} \quad \vdash pcd : \mathbf{pcd} \ t, \pi \quad \vdash pcd' : \mathbf{pcd} \ t', \pi' \quad t'' = t \cap t' \quad \pi'' = \pi \cap \pi'}{\vdash pcd \ \ pcd' : \mathbf{pcd} \ t'', \pi''}$
$\frac{\text{(VAR EXP TYPE)} \quad (var : \mathbf{var} \ t) \in \Pi}{\Pi \vdash var : \mathbf{exp} \ t}$	$\frac{\text{(REGISTER EXP TYPE)} \quad \Pi \vdash e : \mathbf{exp} \ t \quad \text{isClass}(t)}{\Pi \vdash \mathbf{register} \ (e) : \mathbf{exp} \ t}$
$\frac{\text{(CALL EXP TYPE)} \quad \Pi \vdash e : \mathbf{exp} \ c \quad \text{methodBody}(c, m) = t \ m(t_1 \ \mathbf{var}_1, \dots, t_n \ \mathbf{var}_n) \{e\} \quad \Pi \vdash e_1 : \mathbf{exp} \ t_1 \dots \Pi \vdash e_n : \mathbf{exp} \ t_n}{\Pi \vdash e.m(e_1, \dots, e_n) : \mathbf{exp} \ t}$	
$\frac{\text{(EVENT EXP TYPE)} \quad CT(p) = t \ \mathbf{evtype} \ p \ \{t_1 \ \mathbf{var}_1; \dots \ t_n \ \mathbf{var}_n\} \quad \{var_1 : t_1, \dots, var_n : t_n\} \subseteq \Pi \quad \Pi \vdash e : \mathbf{exp} \ t}{\Pi \vdash \mathbf{event} \ p \ \{e\} : \mathbf{exp} \ t}$	
$\frac{\text{(UNDER EXP TYPE)} \quad \Pi \vdash e : \mathbf{exp} \ t}{\Pi \vdash \mathbf{under} \ e : \mathbf{exp} \ t}$	$\frac{\text{(CHAIN EXP TYPE)} \quad \Pi \vdash \mathbf{chain} \ (H, (\mathbf{pcd} \ t, \pi))(e, \rho, \Pi')}{: \mathbf{exp} \ t}$

Figure 11. Type-checking rules for Ptolemy.

$\text{isClass}(t) = (t \in \text{dom}(CT) \wedge CT(t) = \mathbf{class} \ t \dots)$
$\text{isEVType}(t) = (t \in \text{dom}(CT) \wedge CT(t) = t' \ \mathbf{evtype} \ t \dots)$
$\text{isType}(t) = \text{isClass}(t) \vee \text{isEVType}(t)$

Figure 12. Auxiliary functions not in Clifton's dissertation.

The notation $t' \preceq t$ means t' is a subtype of t . It is the reflexive-transitive closure of the declared subclass relationships with the added facts that \top is a supertype of all class types, and that \perp

is a subtype of all class types. The type \perp is used as the type of exceptions. This is formalized in Figure 13.

$$\begin{array}{c}
\text{(BASIS)} \\
\frac{CT(c) = \mathbf{class} \ c \ \mathbf{extends} \ d \{ \dots \}}{c \preceq d} \qquad \text{(REF)} \\
t \preceq t
\end{array}$$

$$\begin{array}{ccc}
\text{(TRANS)} & \text{(TOP)} & \text{(BOTTOM)} \\
\frac{t_1 \preceq t_2 \quad t_2 \preceq t_3}{t_1 \preceq t_3} & \frac{\mathit{isClass}(t)}{t \preceq \top} & \frac{\mathit{isClass}(t)}{\perp \preceq t}
\end{array}$$

Figure 13. Subtyping rules, adapted from [1, Figure 3.4].

1.5 Type Soundness

The proof of soundness of Ptolemy’s type system uses a standard preservation and progress argument [14]. The details are adapted from Clifton’s work [1, 2], which in turn follows Flatt *et al.*’s work [4]. Throughout this section we assume a fixed, well-typed program with a fixed class table.

The key idea in the proof of the subject-reduction theorem is the preservation of consistency between the type environment and the stack and store. This notion is built on the following notion of a (non-null) location having a particular type in the store. This involves fields holding values of their declared types.

DEFINITION 1.1 (*loc* has type *t* in *S*). *Let loc be a location, t be a type, and S be a store. Then loc has type t in S if and only if for some c and F: (i) S(loc) = [c.F], (ii) c ≼ t, (iii) dom(F) = dom(fieldsOf(c)), (iv) rng(F) ⊆ (dom(S) ∪ {null}), and (v) for all f ∈ dom(F), if F(f) = loc', fieldsOf(c)(f) = u, and S(loc') = [c'.F'], then c' ≼ u.*

The key definition of consistency is thus as follows. In the definition, *tenvOf*(ν) is the type environment of a frame ν , and *envOf*(ν) returns ν ’s environment.

DEFINITION 1.2 (Environment-Stack-Store Consistent). *Let Π be a type environment, J a stack, and S a store. Then Π is consistent with (J, S), written $\Pi \approx (J, S)$, if and only if either $J = \bullet$ or $J = \nu + J'$ and all the following hold:*

- $\Pi = \mathit{tenvOf}(\nu)$,
- for $\rho = \mathit{envOf}(\nu)$, and for all $(\mathit{var} : \mathit{var} \ t) \in \Pi$, $\mathit{var} \in \mathit{dom}(\rho)$ and $\rho(\mathit{var})$ has type *t* in *S*, and
- for all $(\mathit{loc} : \mathit{var} \ t) \in \Pi$, $\mathit{loc} \in \mathit{dom}(S)$ and *loc* has type *t* in *S*.

The subject-reduction theorem, as usual, says that evaluation steps preserve both types and consistency. The key idea that makes preservation of consistency easy to prove is the use of type information buried in frames and chain expressions. This type information that is maintained by the operational semantics, but not used by it. Maintenance of this type information occurs each time the stack changes (since the type environment must match that of the top stack frame), and each time a chain expression is created.

THEOREM 1.3 (Subject-reduction). *Let e be an expression, J a stack, S a store, and A an active object list. Let Π be a type environment and t a type. If $\Pi \approx (J, S)$, $\Pi \vdash e : \mathbf{exp} \ t$, and $\langle e, J, S, A \rangle \hookrightarrow \langle e', J', S', A' \rangle$, then there is some Π' and t' such that $\Pi' \vdash e' : \mathbf{exp} \ t'$, $t' \preceq t$ and $\Pi' \approx (J', S')$.*

Proof Sketch: The proof is by cases on the definition of \hookrightarrow (see Figure 6). Assume $\Pi \approx (J, S)$, $\Pi \vdash e : \mathbf{exp} \ t$, and $\langle e, J, S, A \rangle \hookrightarrow \langle e', J', S', A' \rangle$.

The OO cases (rules (NEW), (GET), (SET), (CAST), (NCAST), and (SKIP)) are all straightforward, and can be proved by simple adaptations of Clifton’s proofs for MiniMAO₀ [1, Section 3.1.4]. The result for the exception cases (not shown in Figure 6, but see Figure 8) all follow directly from the use of $\mathbf{exp} \ \perp$ as their type and the fact that the stack in the resulting configuration is empty.

The (CALL) rule is different from Clifton’s MiniMAO₀, and thus must be handled in detail. This case is also a good illustration of how the type information in the configurations is preserved. Suppose $e = \mathit{loc}.m(v_1, \dots, v_n)$. From the hypotheses of the (CALL) rule we have that: $[c.F] = S(\mathit{loc})$, $\nu = \mathbf{execframe} \ m \ \rho \ \Pi'$, $t \ m(t_1 \mathit{var}_1, \dots, t_n \mathit{var}_n) \{e''\} = \mathit{methodBody}(c, m)$, $\rho = \{\mathit{var}_i \mapsto v_i \mid 1 \leq i \leq n\} \oplus (\mathbf{this} \mapsto \mathit{loc})$, and $\Pi' = \{\mathit{var}_i : t_i \mid 1 \leq i \leq n\} \uplus \{\mathbf{this} : c\}$. So in this case, $e' = \mathbf{under} \ e''$, $J' = \nu + J$, and $S' = S$. Since the program is assumed to be well-typed, by the (CHECK PROGRAM) typing rule, all its declarations type check, and so by the (CHECK CLASS) rule, the class *c* type checks, and so by the (CHECK METHOD) rule, the method *m* type checks in class *c*. Thus by the hypotheses of the (CHECK METHOD) rule we can choose Π' to be $\{\mathit{var}_1 : t_1, \dots, \mathit{var}_n : t_n, \mathbf{this} : c\}$. That rule also gives a t' such that $\Pi' \vdash e'' : \mathbf{exp} \ t'$ and $t' \preceq t$. To prove $\Pi' \approx (\nu + J, S')$ we use definition 1.2. The first condition holds by construction, since the type environment in ν is equal to Π' . The second condition holds because for each var_i , if $\rho(\mathit{var}_i) = \mathit{loc}_i \neq \mathbf{null}$, then the loc_i has type t_i in *S*, because for *e* to be well-typed, it must be that $\Pi \vdash v_i : t_i$ (due to the hypotheses of the (CALL EXP TYPE) rule), and by assumption $\Pi \approx (J, S)$. The third condition is vacuous in this case.

The case for the (DEF) rule is similar, and is also similar to Clifton’s (SEQ) case.

Preservation is trivial for the (REGISTER) case, since we can choose $t' = t$. Consistency is also trivial in this case, since the rule makes no changes to the stack or store.

For the (EVENT) rule, suppose $e = \mathbf{event} \ p \ \{e''\}$. From the conclusion of this rule it must be that $J = \nu + J''$ for some ν and J'' . From the hypotheses of the (EVENT) rule we have that: $\rho = \mathit{envOf}(\nu)$, $(t \ \mathbf{evtype} \ p \ \{t_1 \ \mathit{var}_1, \dots, t_n \ \mathit{var}_n\}) \in CT$, $\rho' = \{\mathit{var}_i \mapsto v_i \mid \rho(\mathit{var}_i) = v_i\}$, $\Pi = \mathit{tenvOf}(\nu)$, $\nu' = \mathbf{evframe} \ p \ \rho' \ \pi$, $H = \mathit{hbind}(\nu' + \nu + J, S, A)$, $\pi = \{\mathit{var}_i : t_i \mid 1 \leq i \leq n\}$, and $\theta = \mathbf{pcd} \ t, \pi$. So in this case (substituting the value of θ) $e' = \mathbf{under} \ (\mathbf{chain} \ (H, \mathbf{pcd} \ t, \pi) \ (e'', \rho, \Pi))$, $J' = \nu' + \nu + J''$, and $S' = S$. To preserve consistency, we must choose $\Pi' = \pi$, since that is the type environment of frame ν' . We can also choose $t' = t$, and thus $t' \preceq t$ by reflexivity. With these choices $\pi \vdash e' : \mathbf{exp} \ t$, using the type rules (UNDER EXP TYPE) and (CHAIN EXP TYPE), since e' contains the PCD type $\mathbf{pcd} \ t, \pi$, which determines the type of the chain expression. To prove $\Pi' = \pi \approx (\nu' + \nu + J'', S)$ we use definition 1.2. The first condition holds by construction. The second condition holds because the domain of π is a subset of the domain of Π , π and ρ' are constructed with matching domains, and $S' = S$. The third condition is vacuous in this case.

One of the more interesting cases is the (CHAIN-DONE) rule. Suppose $e = \mathbf{chain} \ (\bullet, \theta) \ (e'', \rho, \Pi')$. From the hypothesis of this rule we have $\nu = \mathbf{contframe} \ \rho \ \Pi''$. So in this case we have $e' = \mathbf{under} \ e''$, $J' = \nu + J$, and $S' = S$. To preserve consistency, we choose $\Pi' = \Pi''$, which is the type environment originally used to type check e'' . Choosing $t' = t$ makes $t' \preceq t$. Then $\Pi' \vdash \mathbf{under} \ e'' : \mathbf{exp} \ t$ using the (UNDER EXP TYPE) rule, and the fact that in chain expressions, e'' and Π'' were placed in the chain expression from an event expression, and the (EVENT EXP TYPE) rule has a hypothesis that shows this. To prove $\Pi' = \Pi'' \approx (\nu + J, S)$ we again use definition 1.2. The first condition holds by construction. The and third conditions hold because the pair of

Π'' and ρ was consistent in the state of the event expression that led to the creation of this chain expression.

For the (CHAIN-RUN) rule, suppose $e = \mathbf{chain} (\langle \langle loc, m, \rho \rangle + H \rangle, \theta) (e'', \rho', \Pi'')$. From the hypothesis of this rule we have $[c.F] = S(loc)$, $\nu = \mathbf{execframe} m \rho'' \Pi'$, $\theta = \mathbf{pcd} t_1, \pi$, $t_1 m(t_1 var_1, \dots, t_n var_n) \{e''\} = \mathit{methodBody}(c, m)$, $n \geq 1$, $\rho'' = \{var_i \mapsto v_i \mid 2 \leq i \leq n, v_i = \rho(var_i)\}$, $\rho''' = \rho'' \oplus \{var_1 \mapsto (\mathbf{chain} (H, \theta) (e'', \rho', \Pi''))\} \oplus \{\mathbf{this} \mapsto loc\}$, $\Pi'' = \{var_i : t_i \mid 1 \leq i \leq n\} \cup \{\mathbf{this} : c\}$. So in this case we have $e' = \mathbf{under} e''$, $J' = \nu + J$, and $S' = S$. To preserve consistency, we choose $\Pi' = \Pi''$. Choosing $t' = t$ makes $t' \preceq t$. Note that from the (CHAIN EXP TYPE) rule, $t = t_1$. Then $\Pi'' \vdash \mathbf{under} e'' : \mathbf{exp} t_1$ using the (UNDER EXP TYPE) rule, and the fact that the method body e'' must have type checked with the rule (CHECK METHOD), since we are assuming that the program type checked. To prove $\Pi' = \Pi'' \approx (\nu + J, S)$ we again use definition 1.2. The first and second conditions hold by construction. In particular, the chain expression used as value that var_1 maps to has type t_1 because θ must be $\mathbf{pcd} t_1, \pi$, due to the way that chain expressions are constructed in the (EVENT) rule and the definition of *match*. The third condition is vacuous in this case. ■

Acknowledgments

Rajan was supported in part by the National Science Foundation grant CNS-0627354. Leavens was supported in part by NSF grant CCF-0429567. The discussions with participants of the seminar course Com S 610-HR, in particular, with Juri Memmert were very helpful.

References

- [1] C. Clifton. A design discipline and language features for modular reasoning in aspect-oriented programs. Technical Report 05-15, Iowa State University, Jul 2005.
- [2] C. Clifton and G. T. Leavens. MiniMAO₁: Investigating the semantics of proceed. *Sci. Comput. Programming*, 63(3):321–374, 2006.
- [3] C. Clifton, G. T. Leavens, and J. Noble. Ownership and effects for more effective reasoning about aspects. In *ECOOP '07, To appear*.
- [4] M. Flatt, S. Krishnamurthi, and M. Felleisen. A programmer's reduction semantics for classes and mixins. In *Formal Syntax and Semantics of Java*, chapter 7, pages 241–269. Springer-Verlag, 1999.
- [5] G. Kiczales et al. An overview of AspectJ. In *ECOOP 2001*, pages 327–353. Springer-Verlag, Jun 2001.
- [6] E. Hilsdale and J. Hugunin. Advice weaving in AspectJ. In *AOSD 04*, pages 26–35, 2004.
- [7] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *OOPSLA '99*, pages 132–146.
- [8] D. C. Luckham and J. Vera. An event-based architecture definition language. *IEEE Trans. Softw. Eng.*, 21(9):717–734, 1995.
- [9] M. Mezini and K. Ostermann. Conquering aspects with Caesar. In *AOSD 03*, pages 90–99.
- [10] G. Nadathur and D. J. Mitchell. System description: Teyjus - a compiler and abstract machine based implementation of lambda-prolog. In *CADE-16: Proceedings of the 16th International Conference on Automated Deduction*, pages 287–291, 1999.
- [11] H. Rajan and K. J. Sullivan. Classpects: unifying aspect- and object-oriented language design. In *ICSE 2005*, pages 59–68.
- [12] H. Rajan and K. J. Sullivan. Eos: instance-level aspects for integrated system design. In *ESEC/FSE 2003*, pages 297–306, Sep. 2003.
- [13] D. A. Schmidt. *The Structure of Typed Programming Languages*. Foundations of Computing Series. MIT Press, MA, 1994.
- [14] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, Nov 1994.