

Sequence Traces for Object-Oriented Executions

Carl Eastlund
Northeastern University
Boston, Massachusetts
cce@ccs.neu.edu

Matthias Felleisen
Northeastern University
Boston, Massachusetts
matthias@ccs.neu.edu

ABSTRACT

Over the past 30 years, researchers have developed a large variety of semantic models of object-oriented computations. These include object calculi as well as denotational, small-step operational, big-step operational, and reduction semantics. None of them, however, separate the object-oriented aspects of a computation from the procedural or applicative ones; none express the object-oriented computations via message-based communication among pools of objects.

In this paper, we present a novel, two-level framework of object-oriented computation. Our purpose is to create a foundation for program inspection and debugging tools that enable programmers to focus on the object-oriented part of a computation and to zoom in on the internal actions of objects only as needed. The upper level of the framework borrows elements from UML's sequence diagrams to express the message exchanges among objects. The lower level is a parameter of the upper level; it models all those elements of a programming language that are procedural. We include several instantiations of the framework as well as its application to a uniform (parameterized) type soundness proof.

1. MODELS OF EXECUTION

Some 30 years ago, Hewitt [20, 21] introduced the ACTOR model of computation, which is arguably the first model of object-oriented computation. Since then, people have explored a range of mathematical models of object-oriented program execution: denotational semantics of objects and classes [7, 8, 22, 30]; object calculi [1]; small step and big step operational semantics [9]; reduction semantics [13]; formal variants of ACTOR [2]; and so on [4, 18].

While all of these semantic models have their appropriate uses and contributed to the community's understanding of object-oriented languages, none of them isolate the intuition of pools of objects that exchange messages with each other. As the messages go back and forth, objects change state; get to know more about other objects; and the pool of objects grows. Everything else is object-internal.

In contrast to semantic models of classes and objects, specification frameworks such as UML [14] provide graphical languages for describing classes, objects, and their desired behaviors. In particular, UML provides (among other notations) *sequence diagrams*. Requirements analysts and specifiers use them to record the desired time-line of interactions among (classes of) objects, repetitions, constraints, and other expected details.

Ideally, language researchers should provide a semantic framework that matches such object-oriented specifications of executions. Tools based on such a framework would then present programmers a view of the computation that is related to their specifications. When needed, they can expose internal actions of methods and objects. In this paper, we introduce such a two-level framework. The *upper level* represents all object-oriented actions during a program's execution. It tracks six kinds of actions: object creation, field inspection, field mutation, object tag inspection, method calls and method return; we do not consider any other action an object-oriented computation. Computations at this upper level are (equivalent to) manipulations of UML-like sequence diagrams; to assign a precise meaning to these manipulations, each diagram is expressible as a configuration and the manipulations are formulated as rewriting actions on configurations [23].

The upper level is parameterized over the internal semantics of method bodies, dubbed the *lower level*. Technically the parameterization takes the form of several sets and functions. To instantiate the framework for a specific language, a semanticist must map the object-oriented part of a language to the object-oriented level of the framework and must express the remaining actions as the lower level. Thus, the lower level can be represented as a state machine, a set of mathematical functions, or whatever else a semanticist finds appropriate. We demonstrate how to instantiate the framework with several examples.

In addition to developing a precise mathematical meaning for the framework, we have also implemented a prototype of the framework. The prototype is a component of DrScheme [11] for the kernel of PLT Scheme's class system [12]. It draws a UML-like trace of the program's object-oriented actions and allows programmers to inspect object states.

The rest of the paper is organized as follows. Section 2 presents a high-level overview of the prototype, motivating the development of the semantics. Section 3 introduces the framework and section 4 demonstrates how to instantiate it for three radically different languages. Section 5 shows how

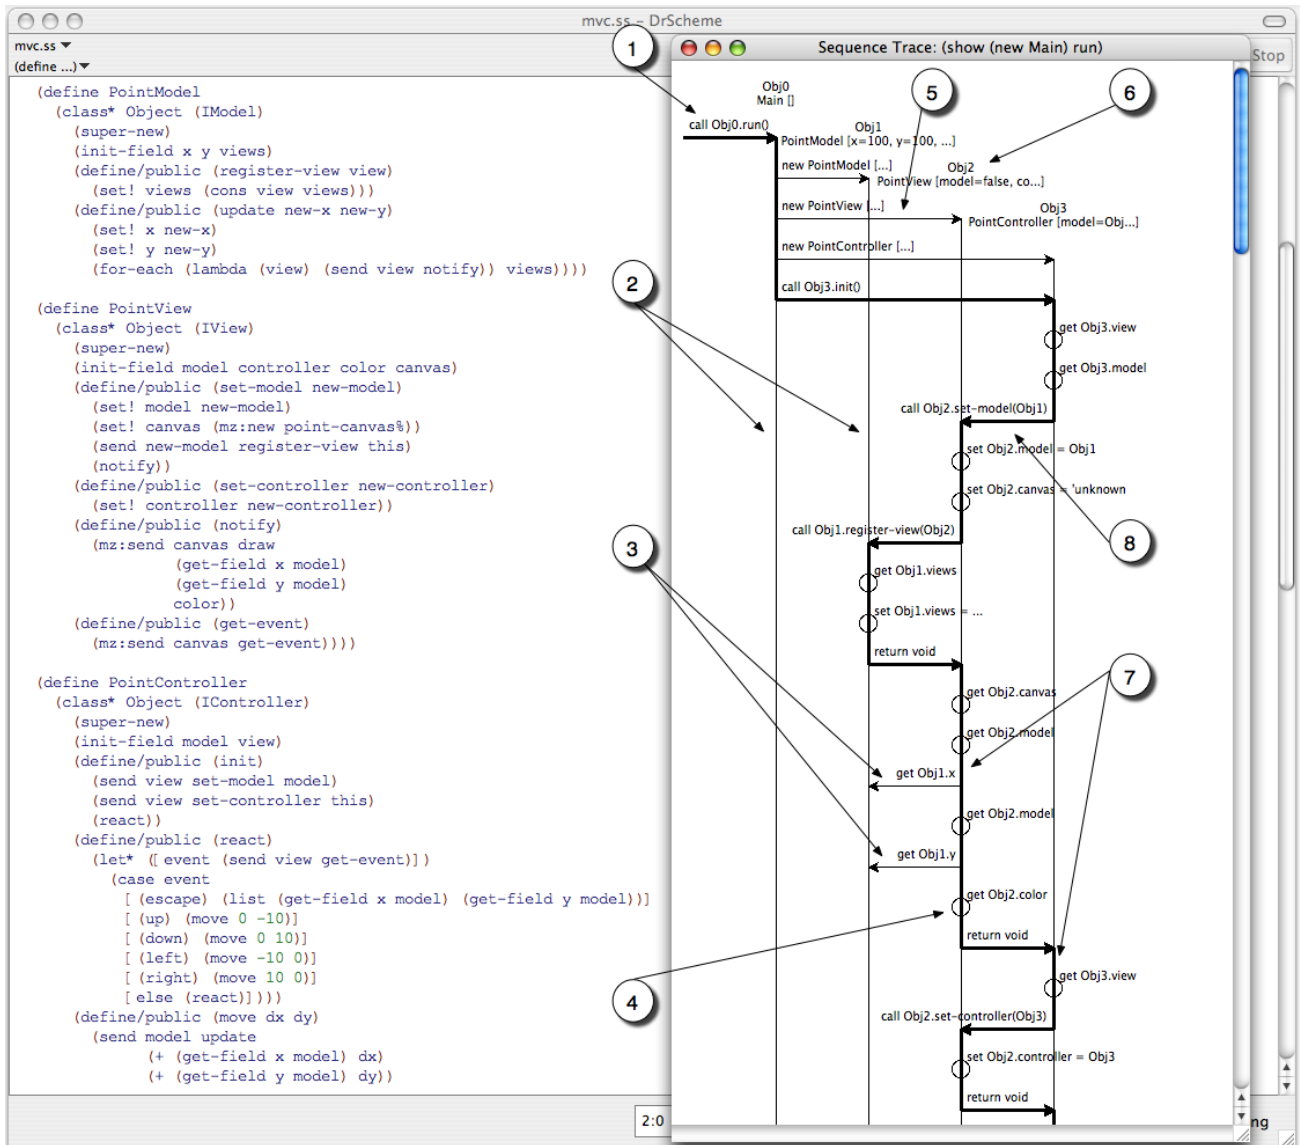


Figure 1: Graphical sequence trace of a Model-View-Controller program.

to prove most of the type soundness theorem at the upper level and applies this to one of the models from section 4. The last two sections are about related and future work.

What we do *not* do: The purpose of this paper is *not* to assign meaning to UML. We simply borrow some of UML’s graphical elements. The following three sections clarify this difference implicitly; we explain it in some detail in the section on related work.

2. SEQUENCE TRACES

Good software development requires good software tools, including tools for inspecting program executions and debugging code. While many such tools exist, most are generic and don’t capture the essence of object-oriented programming. In particular, debuggers often expose an execution state through stack traces and variables [25], which are procedural concepts rather than object-oriented ones. Some trace tools record events without an object-oriented model; others translate traces into abstract interpretations [10, 31, 34].

We have prototyped a tool that presents the execution of object-oriented programs in the style of UML sequence diagrams. The tool operates on the PLT Scheme class system [12]. Roughly speaking, the object-oriented portion of PLT Scheme is an untyped version of Java’s class and interface system, extended with mixins in the spirit of Flatt et al [13]. Objects interact through fields, methods, constructors, and reflection. The tool captures a *sequence trace* of these interactions, reconstructing a snapshot of the program state at each object-oriented action.

Figure 1 depicts a sample sequence trace. The pictured program excerpt (on the left) is an instance of the Model-View-Controller pattern written with PLT Scheme classes. The visible portion of the trace (on the right) is a prefix of the program’s execution, starting with the invocation of the `Main` object (number 1 on the figure). Following UML tradition, the vertical lines are object lifelines (number 2) representing the existence of an object. Horizontal arrows are messages (number 3) sent from one object to another; circles represent messages sent from an object to itself (number 4). New lifelines start after `new` messages (number 5). Every lifeline has a header (number 6) with a name for the object above and its class and fields below. The control flow of the execution is emphasized by the bold segments of messages and lifelines. Where an object lifeline is highlighted (number 7), that object has control. Where a message arrow is highlighted (number 8), one object transfers control to another.

Sequence traces, such as the one in Figure 1, suggest a model of computation as communication. In this model, an execution for an object-oriented program is represented as a collection of object lifelines and the messages passed between them. The model “hides” computations that take place inside of methods and that don’t have any externally visible communication. Of course, this model is not particular to PLT Scheme; it is the core of any object-oriented programming language and deserves a formal exploration.

3. FRAMEWORK SEMANTICS

Our framework assigns semantics to object-oriented programming languages at two levels. The upper level is all about objects, their creation, their life lines, and their ex-

| | |
|-------------------------------|--|
| NAME^L | syntactic category NAME in language L , or a metavariable in it (by context) |
| \vec{t} | any number of elements of the form t |
| $c[e]$ | expression e in evaluation context c |
| $e[v/x]$ | substitution of v for free x in expression e |
| $d \xrightarrow{p} r$ | partial functions from domain d to range r |
| $d \xrightarrow{f} r$ | finite mappings from domain d to range r |
| $\overline{[a \mapsto b]}$ | finite mapping of each a to b |
| $f[\overline{[a \mapsto b]}]$ | finite mapping extension by juxtaposition overrides mappings in f |

Figure 2: Notational conventions.

changes of messages. The lower level concerns all those aspects of a language’s semantics that are unrelated to its object-oriented nature, e.g., static methods, blocks, decision constructs, looping constructs, and so on.

Figure 3 gives the full syntax of the upper level using the notation in Figure 2. The upper level represents the execution of an object-oriented program via a TRACE, which is a sequence of STATES. Each state is a four-tuple, consisting of a POOL of objects, a STACK, an object reference, and an ACTION. A pool is a finite map from object references to OBJECTS. The latter have two parts: a DYNAMIC record, mapping FIELDS to VALUES, and a STATIC record, containing all invariant parts of an object. The stack is a sequence of object references and so-called CONTINUATIONS.

Finally, there are six kinds of messages: a `new` message, which contains an OBJECT to allocate; an `inspect` message for a designated object reference; a `get` message for a specific field in a designated object reference; a `set` message with a value for a specific field in a designated object reference; a `call` message for a specific method with a number of argument values; and a `return` message with a value. Most messages contain a continuation (using notation \triangleright). When an object is unable to send a message, it may signal an error instead.

In order to accommodate as many object-oriented languages as possible, this syntactic part of the upper level is parameterized over a number of language-specific sets: the text of PROGRAMS, FIELDS (field names); METHODS (method names); PRIMITIVE non-object values; STATIC; CONTS, which represent method-local continuations; and LANGERROR, the collection of language-specific errors.

Figure 4 gives the upper level operational semantics of sequence traces, along with descriptions and signatures for its lower level parameters. A trace is the result of rewriting the initial state, step by step, into a final state. The parameter *init* is a function mapping a program to its initial state. Each subsequent state depends on the preceding action, as follows:

object creation A `new` action adds a reference and an object to the pool. The initiating object retains control.

object inspection An `inspect` action retrieves the static record of an object reference from the pool.

field lookup A `get` action retrieves the value of a field from an object in the pool.

field update A `set` action changes the value of a field in an object in the pool.

| | | |
|-----------|--|--|
| PROG | lower level definition | Program |
| CONT | lower level definition | Continuation |
| STATIC | lower level definition | Static Record |
| FIELD | lower level definition | Field Name |
| METHOD | lower level definition | Method Name |
| PRIM | lower level definition | Primitive Value |
| LANGERROR | lower level definition | Language-specific Error |
| REF | countable set | Object Reference |
| TRACE | = $\overrightarrow{\text{STATE}}$ | Sequence Trace |
| STATE | = $\langle \text{POOL}, \text{STACK}, \text{REF}, \text{ACTION} \rangle$ | Execution State |
| POOL | : $\text{REF} \xrightarrow{f} \text{OBJECT}$ | Object Pool |
| STACK | = $\epsilon \mid \langle \text{REF}, \text{CONT} \rangle \text{STACK}$ | Method Stack |
| OBJECT | = STATIC DYNAMIC | Object Record |
| DYNAMIC | : $\text{FIELD} \xrightarrow{f} \text{VALUE}$ | Dynamic Record |
| VALUE | = $\text{PRIM} \mid \text{REF} \mid \text{STATIC}$ | Value |
| ACTION | = $\text{MESSAGE} \mid \text{ERROR}$ | Object-Oriented Action |
| MESSAGE | = new OBJECT \triangleright CONT inspect REF \triangleright CONT get REF.FIELD \triangleright CONT set REF.FIELD := VALUE \triangleright CONT call REF.METHOD($\overrightarrow{\text{VALUE}}$) \triangleright CONT return VALUE | Object Creation Object Inspection Field Lookup Field Update Method Call Method Return |
| RESULT | = $\langle \text{POOL}, \epsilon, \text{REF}, \text{return VALUE} \rangle$ $\langle \text{POOL}, \text{STACK}, \text{REF}, \text{ERROR} \rangle$ | Success Result Failure Result |
| ERROR | = LANGERROR error:reference error:field | Execution Error |

Figure 3: Upper-level definition of sequence trace syntax.

method call A **call** action invokes a method in an object on a number of arguments.

method return A **return** action completes the current method call.

The **new** and **set** messages update the object pool; **new** adds new objects and **set** replaces their field values. The **call** and **return** messages manipulate the stack and the reference to the controlling object, transferring control between objects. The former pushes a frame onto the stack containing the current object reference and the action’s continuation; the latter removes and reinstates them later. Aside from these cases, messages leave the pool, stack, and controlling object reference unchanged.

Every message changes the action field. At each step, the rewriting system uses either the function *invoke* or *resume* to compute the next action. Both functions are parameters of the rewriting system. The former (*invoke*) maps a method call for a designated object, method name, and a sequence of arguments to an action; the latter maps a value and a continuation to an action. Both functions are partial, admitting the possibility of non-termination. Also, both functions may map their inputs to a language-specific error.

3.1 To Diagrams and Back

The formal definition of sequence traces represents their visual form, but does not resemble it. Still, it is possible to convert one form into the other (relative to internal computation). The example in Figure 5 shows a visual sequence trace next to its formal representation.

Conversion from a formal trace to a visual one proceeds state by state “down” the execution. To draw a visual sequence trace from a formal one, we start with the initial

object pool: a lifeline header for each object in the pool. A lifeline header is an object’s reference above its static and dynamic records. For each state in the trace, we then perform the following steps:

1. Extend the lifeline of all objects in the pool.
2. If an object is created by a **new** action, add a new object header to the right of the object pool
3. Draw an arrow for the current action. The arrow starts at the controlling object, denoted by the reference in the current state. Most messages refer to a target object, where the arrow ends. A **new** message’s arrow ends at the new object header. A **return** message’s arrow ends at the object on top of the stack.
4. Label the arrow with the current action. If the action has a continuation, leave it off. Continuations represent internal actions of an object and are not part of a visual sequence trace.

To construct a formal sequence trace from a picture, we need to accumulate the state of the program along the trace. The initial store is simply the collection of lifeline headers at the top of the trace. The initial stack is empty; the controlling object is the source of the first arrow and the action is the arrow’s label. Any continuations must be inferred since they are not shown in visual sequence traces. For the rest of the diagram, we use these steps.

1. After a **new** message, we add the new object header to the object pool. After a **set** message, we replace the appropriate field value. The object pool remains otherwise unchanged.

Lower level functions

| | | |
|-----------------|--|---------------------------------------|
| $init$ | $PROG \longrightarrow STATE$ | Constructs the initial program state. |
| $invoke_{PROG}$ | $\langle REF, OBJECT, METHOD, \overline{VALUE} \rangle \xrightarrow{P} ACTION$ | Invokes a method. |
| $resume_{PROG}$ | $\langle CONT, VALUE \rangle \xrightarrow{P} ACTION$ | Resumes a suspended computation. |

Evaluation

$PROG \downarrow RESULT$ means $init(PROG) \mapsto_{PROG}^* RESULT$

where

Object Creation

$\langle POOL, STACK, REF, \mathbf{new} \text{ OBJECT} \triangleright CONT \rangle$
 $\mapsto_{PROG} \langle POOL[REF' \mapsto OBJECT], STACK, REF, resume_{PROG}(CONT, REF') \rangle$
 where $REF' \notin dom(POOL)$

Object Inspection

$\langle POOL, STACK, REF, \mathbf{inspect} \text{ REF}' \triangleright CONT \rangle$
 $\mapsto_{PROG} \langle POOL, STACK, REF, resume_{PROG}(CONT, STATIC) \rangle$
 if $POOL(REF') = STATIC \text{ DYNAMIC}$

Field Lookup

$\langle POOL, STACK, REF, \mathbf{get} \text{ REF}' \cdot \text{FIELD} \triangleright CONT \rangle$
 $\mapsto_{PROG} \langle POOL, STACK, REF, resume_{PROG}(CONT, VALUE) \rangle$
 if $POOL(REF') = STATIC \text{ DYNAMIC}$ and $DYNAMIC(FIELD) = VALUE$

Field Update

$\langle POOL, STACK, REF, \mathbf{set} \text{ REF}' \cdot \text{FIELD} := VALUE \triangleright CONT \rangle$
 $\mapsto_{PROG} \langle POOL[REF' \mapsto STATIC \text{ DYNAMIC}[FIELD \mapsto VALUE]], STACK, REF, resume_{PROG}(CONT, VALUE) \rangle$
 if $POOL(REF') = STATIC \text{ DYNAMIC}$ and $FIELD \in dom(DYNAMIC)$

Method Call

$\langle POOL, STACK, REF, \mathbf{call} \text{ REF}' \cdot \text{METHOD}(\overline{VALUE}) \triangleright CONT \rangle$
 $\mapsto_{PROG} \langle POOL, \langle REF, CONT \rangle \text{ STACK}, REF', invoke_{PROG}(REF', POOL(REF'), METHOD, \overline{VALUE}) \rangle$
 if $REF' \in dom(POOL)$

Method Return

$\langle POOL, \langle REF', CONT \rangle \text{ STACK}, REF, \mathbf{return} \text{ VALUE} \rangle$
 $\mapsto_{PROG} \langle POOL, STACK, REF', resume_{PROG}(CONT, VALUE) \rangle$

Reference Error

$\langle POOL, STACK, REF, MESSAGE \rangle$
 $\mapsto_{PROG} \langle POOL, STACK, REF, \mathbf{error:reference} \rangle$
 if $MESSAGE = \mathbf{inspect} \text{ REF}' \triangleright CONT$
 | $\mathbf{get} \text{ REF}' \cdot \text{FIELD} \triangleright CONT$
 | $\mathbf{set} \text{ REF}' \cdot \text{FIELD} := VALUE \triangleright CONT$
 | $\mathbf{call} \text{ REF}' \cdot \text{METHOD}(\overline{VALUE}) \triangleright CONT$
 and $REF' \notin dom(POOL)$

Field Error

$\langle POOL, STACK, REF, MESSAGE \rangle$
 $\mapsto_{PROG} \langle POOL, STACK, REF, \mathbf{error:field} \rangle$
 if $MESSAGE = \mathbf{get} \text{ REF}' \cdot \text{FIELD} \triangleright CONT$ | $\mathbf{set} \text{ REF}' \cdot \text{FIELD} := VALUE \triangleright CONT$
 and $POOL(REF') = STATIC \text{ DYNAMIC}$ and $FIELD \notin dom(DYNAMIC)$

Figure 4: Upper-level definition of sequence trace semantics.

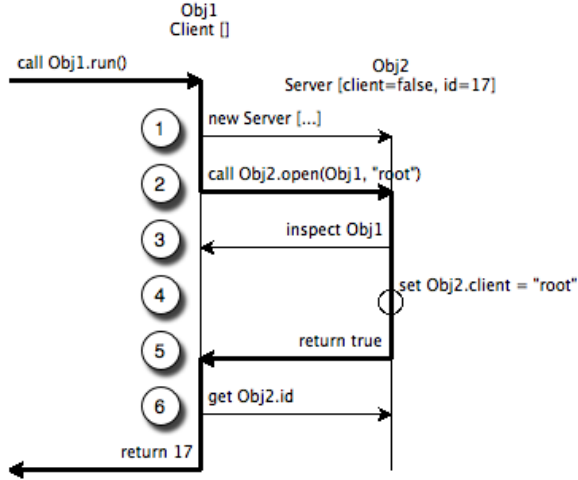


Figure 5: Correlation between visual and formal sequence traces.

2. After a **call** message, we add the current object reference to the stack. We install the message’s receiver as the new object reference. After a **return** message, we remove a reference from the top of the stack and install it as the new object reference. The stack and object reference are otherwise unchanged.
3. The new state’s action is always the label on the arrow in the next row. The action’s continuation must be inferred.

4. FRAMEWORK INSTANTIATIONS

The framework accommodates a variety of programming language models. To demonstrate its flexibility in this regard, we demonstrate how to formulate three radically different, well-known models as lower levels of our framework: Abadi and Cardelli’s object calculus [1] in section 4.1, Flatt et alii’s ClassicJava [13] in section 4.2, and Castagna et alii’s λ -calculus [5] in section 4.3. The first is a pure calculus of objects; the second is a semantics of a class-based object-oriented language in the spirit of Java, and the last is a model used to formalize multiple dispatch.

Formulating a language semantics in our framework has two purposes. Conceptually, it helps a semanticist determine what is object-oriented about a language and what is procedural. This point comes out especially well with our reformulation of ClassicJava and the modeling of super-method calls. It is strictly impossible in conventional models to discuss this issue. Pragmatically, the formulation of a language in our framework is the first step for instantiating our tracing tool for this language. That is, once we have a framework-based semantics, it is straightforward to instrument the compiler and to combine it with our tracing component to get a new debugging tool.

4.1 Object Calculus Sequence Traces

Abadi and Cardelli’s object calculus [1] is the quintessential model of object-oriented computation; its role is similar to that of the λ -calculus in functional programming. The purpose of the calculus is to model objects and methods;

1. $\langle \text{POOL}_1, \epsilon, \text{Obj1}, \text{new Server DYNAMIC}_1 \triangleright \text{CONT}_1 \rangle$
where $\text{POOL}_1 = [\text{Obj1} \mapsto \text{Client } []]$
and $\text{DYNAMIC}_1 = [\text{client} \mapsto \text{false}, \text{id} \mapsto 17]$
2. $\langle \text{POOL}_2, \epsilon, \text{Obj1}, \text{call Obj2.open}(\text{Obj1}, \text{"root"}) \triangleright \text{CONT}_2 \rangle$
where $\text{POOL}_2 = \text{POOL}_1[\text{Obj2} \mapsto \text{Server DYNAMIC}_1]$
3. $\langle \text{POOL}_2, \langle \text{Obj1}, \text{CONT}_2 \rangle, \text{Obj2}, \text{inspect Obj1} \triangleright \text{CONT}_3 \rangle$
4. $\langle \text{POOL}_2, \langle \text{Obj1}, \text{CONT}_2 \rangle, \text{Obj2}, \text{set Obj2.client := "root"} \triangleright \text{CONT}_4 \rangle$
5. $\langle \text{POOL}_3, \langle \text{Obj1}, \text{CONT}_2 \rangle, \text{Obj2}, \text{return true} \rangle$
where $\text{POOL}_3 = \text{POOL}_2[\text{Obj2} \mapsto \text{Server DYNAMIC}_2]$
and $\text{DYNAMIC}_2 = \text{DYNAMIC}_1[\text{client} \mapsto \text{"root"}]$
6. $\langle \text{POOL}_3, \epsilon, \text{Obj1}, \text{get Obj2.id} \triangleright \text{CONT}_5 \rangle$

in particular, objects don’t have fields, only constant-value methods. The calculus has three actions: object creation, method invocation, and method override. Object creation specifies the names and bodies of methods. Method invocation specifies an object (the receiver) and a method name; methods have no external arguments. Method override provides an object, a method name, and the new body for that method. This creates a new object with the updated method body, inheriting all other methods from the original object.

This model is a natural fit for our framework. The bodies of all methods become the static record of an object. Object creation and method invocation are already framework messages. Method override takes two messages: see below.

Figure 6 shows the lower level parameters defining object calculus syntax. Programs in the object calculus are expressions. Continuations are evaluation contexts over the expressions. An expression can be a variable, reference, method call, method override, or object constructor. An object’s static record maps its method names to method bodies. Finally, the lower level also defines method errors.

Figure 7 defines the semantics of the object calculus. Program initialization stores the initial expression in a method body and invokes that method as the first action. The *resume* and *invoke* functions are defined in terms of an *eval* function, operating on expressions via a reduction relation.

Object creation and method invocation evaluate directly to the appropriate messages; all method bodies are stored in an object’s static record, and extracted when invoked. Method override is modelled as two messages. The first extracts the static record of the receiver to copy the inherited methods. The second constructs the updated object.

4.2 ClassicJava Sequence Traces

ClassicJava [13] is another widely used model of object-oriented computation. In contrast to the object calculus, it is class-based, including most of Java’s core features such as classes, interfaces, inheritance, method overriding, and typecasts. Naturally, this does not match our framework perfectly, but the transition is remarkably smooth.

Figure 8 shows the lower-level parameters defining Clas-

| | | | |
|--------------------------|---|---|--------------------------|
| PROG^ζ | = | EXPR^ζ | Object Calculus Program |
| CONT^ζ | = | \square | Empty Context |
| | | $\text{CONT}^\zeta.\text{METHOD}^\zeta$ | Method Call Receiver |
| | | $\text{CONT}^\zeta.\text{METHOD}^\zeta \leftarrow \text{BODY}^\zeta$ | Method Override Receiver |
| STATIC^ζ | : | $\text{METHOD}^\zeta \xrightarrow{f} \text{BODY}^\zeta$ | Static Record |
| FIELD^ζ | = | \emptyset | Field Names |
| METHOD^ζ | = | method label | Method Names |
| PRIM^ζ | = | \emptyset | Primitive Values |
| LANGERROR^ζ | = | error:method | Method Error |
| where | | | |
| BODY^ζ | = | $\zeta(\text{VAR}).\text{EXPR}^\zeta$ | Method Body |
| EXPR^ζ | = | $\text{VALUE} \mid \text{VAR}$ | Simple Expressions |
| | | $\text{EXPR}^\zeta.\text{METHOD}^\zeta$ | Method Call |
| | | $\text{EXPR}^\zeta.\text{METHOD}^\zeta \leftarrow \text{BODY}^\zeta$ | Method Override |
| | | $\overrightarrow{\{\text{METHOD}^\zeta = \text{BODY}^\zeta\}}$ where $(\text{METHOD}_i^\zeta \text{ distinct})$ | Object Construction |

Figure 6: Lower-level parameters for object calculus syntax.

| | | |
|--|---|---|
| $\text{init}^\zeta(\text{EXPR})$ | = | $\langle [\text{REF}_0 \mapsto [\text{METHOD}_0 \mapsto \zeta(\text{VAR}_0).\text{EXPR}]],$ $\epsilon, \text{REF}_0, \text{call } \text{REF}_0.\text{METHOD}_0() \triangleright \square \rangle$ |
| $\text{resume}_{\text{PROG}}^\zeta(\text{CONT}, \text{VALUE})$ | = | $\text{eval}^\zeta(\text{CONT}[\text{VALUE}])$ |
| $\text{invoke}_{\text{PROG}}^\zeta(\text{REF}, \text{STATIC}, \text{METHOD}, \langle \rangle)$ | = | $\text{eval}^\zeta(\text{EXPR}[\text{REF}/\text{VAR}])$ where $\text{STATIC}(\text{METHOD}) = \zeta(\text{VAR}).\text{EXPR}$ |
| $\text{invoke}_{\text{PROG}}^\zeta(\text{REF}, \text{STATIC}, \text{METHOD}, \langle \rangle)$ | = | error:method where $\text{METHOD} \notin \text{dom}(\text{STATIC})$ |
| $\text{invoke}_{\text{PROG}}^\zeta(\text{REF}, \text{STATIC}, \text{METHOD}, \overrightarrow{\text{VALUE}})$ | = | error:arity where $ \overrightarrow{\text{VALUE}} > 0$ |
| • $\text{eval}^\zeta : \text{EXPR} \longrightarrow \text{ACTION}$ is: | | |
| $\text{eval}^\zeta(\text{VALUE})$ | = | return VALUE |
| $\text{eval}^\zeta(\text{CONT}[\overrightarrow{\{\text{METHOD} = \text{BODY}\}}])$ | = | new $[\text{METHOD} \mapsto \text{BODY}] \triangleright \text{CONT}$ |
| $\text{eval}^\zeta(\text{CONT}[\text{REF}.\text{METHOD}])$ | = | call $\text{REF}.\text{METHOD}() \triangleright \text{CONT}$ |
| $\text{eval}^\zeta(\text{CONT}[\text{REF}.\text{METHOD} \leftarrow \text{BODY}])$ | = | inspect $\text{REF} \triangleright \text{CONT}[\square].\text{METHOD} \leftarrow \text{BODY}$ |
| $\text{eval}^\zeta(\text{CONT}[\text{STATIC}.\text{METHOD} \leftarrow \text{BODY}])$ | = | new $\text{STATIC}[\text{METHOD} \mapsto \text{BODY}] \triangleright \text{CONT}$ where $\text{METHOD} \in \text{dom}(\text{STATIC})$ |
| $\text{eval}^\zeta(\text{CONT}[\text{STATIC}.\text{METHOD} \leftarrow \text{BODY}])$ | = | error:method where $\text{METHOD} \notin \text{dom}(\text{STATIC})$ |

Figure 7: Lower-level parameters for object calculus semantics.

| | | |
|------------------|---|---------------------------|
| DEF^{CJ} | class or interface definition | Definition |
| $CLASS^{CJ}$ | class name | Class |
| $TNAME^{CJ}$ | class or interface name | Type |
| $FIELD^\ell$ | countable set | Field Label |
| $METHOD^\ell$ | countable set | Method Label |
| $PROG^{CJ}$ | $= \overrightarrow{DEF^{CJ}} \text{ new } CLASS^{CJ}.METHOD^\ell()$ | Program |
| $CONT^{CJ}$ | $= []$ | Empty Context |
| | $CONT^{CJ} : CLASS^{CJ}.FIELD^\ell$ | Field Lookup Receiver |
| | $CONT^{CJ} : CLASS^{CJ}.FIELD^\ell = EXPR^{CJ}$ | Field Update Receiver |
| | $VALUE : CLASS^{CJ}.FIELD^\ell = CONT^{CJ}$ | Field Update Value |
| | $CONT^{CJ}.METHOD^\ell(\overrightarrow{EXPR^{CJ}})$ | Method Call Receiver |
| | $VALUE.METHOD(\overrightarrow{VALUE} \overrightarrow{CONT^{CJ}} \overrightarrow{EXPR^{CJ}})$ | Method Call Argument |
| | super = $REF : CLASS^{CJ}.METHOD^\ell(\overrightarrow{VALUE} \overrightarrow{CONT^{CJ}} \overrightarrow{EXPR^{CJ}})$ | Super Call Argument |
| | view $TNAME \overrightarrow{CONT^{CJ}}$ | Typecast Argument |
| | view $CONT^{CJ} \leq TNAME \ VALUE$ | Elaborated Typecast Class |
| | let $VAR = CONT^{CJ} \text{ in } EXPR^{CJ}$ | Let Binding R.H.S. |
| $STATIC^{CJ}$ | $= CLASS^{CJ}$ | Static Record |
| $FIELD^{CJ}$ | $= \langle CLASS^{CJ}, FIELD^\ell \rangle$ | Field Name |
| $METHOD^{CJ}$ | $= METHOD^\ell \mid \langle CLASS^{CJ}, METHOD^\ell \rangle$ | Method Name |
| $PRIM^{CJ}$ | $= \text{null}$ | Primitive Value |
| $LANGERROR^{CJ}$ | $= \text{error:method}$ | Method Error |
| | error:null | Null Error |
| | error:typecast | Typecast Error |
| where | | |
| $EXPR^{CJ}$ | $= REF$ | Object Reference |
| | new $CLASS^{CJ}$ | Object Creation |
| | VAR | Variable |
| | this | Self Reference |
| | null | Null Value |
| | $EXPR^{CJ} : CLASS.FIELD$ | Field Lookup |
| | $EXPR^{CJ} : CLASS.FIELD = EXPR^{CJ}$ | Field Update |
| | $EXPR^{CJ}.METHOD(\overrightarrow{EXPR^{CJ}})$ | Method Call |
| | super = this : $CLASS.METHOD(\overrightarrow{EXPR^{CJ}})$ | Super Call |
| | view $TNAME \ EXPR^{CJ}$ | Typecast |
| | view $CLASS \leq TNAME \ EXPR^{CJ}$ | Elaborated Typecast |
| | let $VAR = EXPR^{CJ} \text{ in } EXPR^{CJ}$ | Let Binding |

Figure 8: Lower level parameters for ClassicJava syntax.

| | | |
|--|--------|--|
| $TNAME \leq_{PROG}^{CJ} TNAME'$ | \iff | $TNAME$ is a subtype of $TNAME'$ |
| $\langle CLASS', FIELD, TNAME \rangle \in_{PROG}^{CJ} CLASS$ | \iff | class $CLASS$ includes or inherits field $FIELD$ of type $TNAME$ defined in class $CLASS'$ |
| $\langle METHOD, \overrightarrow{TNAME} \rightarrow TNAME, \overrightarrow{VAR}, EXPR \rangle \in_{PROG}^{CJ} CLASS$ | \iff | class $CLASS$ includes or inherits a definition for method $METHOD$ with type $\overrightarrow{TNAME} \rightarrow TNAME$, input variables \overrightarrow{VAR} , and method body $EXPR$. |

Figure 9: ClassicJava semantic relations.

$$\begin{aligned}
init^{CJ}(\overrightarrow{\text{DEF new CLASS.METHOD()}}) &= \langle [REF_0 \mapsto \text{newobject}^{CJ}(\text{CLASS})], \\
&\quad \epsilon, REF_0, \text{call } REF_0.METHOD() \triangleright [] \rangle \\
resume_{PROG}^{CJ}(\text{CONT}, \text{VALUE}) &= eval_{PROG}^{CJ}(\text{CONT}[\text{VALUE}]) \\
invoke_{PROG}^{CJ}(\text{REF}, \text{CLASS DYNAMIC}, \text{METHOD}, \overrightarrow{\text{VALUE}}) &= evalmethod_{PROG}^{CJ}(\text{REF}, \text{CLASS}, \text{METHOD}, \overrightarrow{\text{VALUE}}) \\
invoke_{PROG}^{CJ}(\text{REF}, \text{OBJECT}, \langle \text{CLASS}, \text{METHOD} \rangle, \overrightarrow{\text{VALUE}}) &= evalmethod_{PROG}^{CJ}(\text{REF}, \text{CLASS}, \text{METHOD}, \overrightarrow{\text{VALUE}})
\end{aligned}$$

• $\text{newobject}^{CJ} : \text{CLASS} \rightarrow \text{OBJECT}$ is:

$$\begin{aligned}
\text{newobject}^{CJ}(\text{CLASS}) &= \text{CLASS} [\langle \text{CLASS}', \text{FIELD} \rangle \mapsto \text{null}] \\
&\quad \text{where } \langle \text{CLASS}', \text{FIELD} \rangle = \{ \langle \text{CLASS}', \text{FIELD} \rangle \mid \langle \text{CLASS}', \text{FIELD}, \text{TNAME} \rangle \in_{PROG}^{CJ} \text{CLASS} \}
\end{aligned}$$

• $\text{evalmethod}^{CJ} : \langle \text{REF}, \text{CLASS}, \text{METHOD}, \overrightarrow{\text{VALUE}} \rangle \rightarrow \text{ACTION}$ is:
 $\text{evalmethod}_{PROG}^{CJ}(\text{REF}, \text{CLASS}, \text{METHOD}, \overrightarrow{\text{VALUE}}) = \text{ACTION}$ by the following table:

| ACTION | where |
|--|--|
| $eval_{PROG}^{CJ}(\text{EXPR}[\text{REF}/\text{this}][\overrightarrow{\text{VALUE}}/\overrightarrow{\text{VAR}}])$ | $\langle \text{METHOD}, \text{TNAME} \rightarrow \text{TNAME}, \overrightarrow{\text{VAR}}, \text{EXPR} \rangle \in_{PROG}^{CJ} \text{CLASS}$ and $ \overrightarrow{\text{VALUE}} = \overrightarrow{\text{VAR}} $ |
| error:method | otherwise |

• $eval^{CJ} : \text{EXPR} \rightarrow \text{ACTION}$ is:

$eval_{PROG}^{CJ}(\text{EXPR}) = \text{ACTION}$ where $\text{EXPR} \mapsto_{PROG}^* \text{EXPR}'$ by the following table:

| ACTION | EXPR' |
|--|--|
| return VALUE | VALUE |
| new $\text{newobject}^{CJ}(\text{CLASS}) \triangleright \text{CONT}$ | $\text{CONT}[\text{new CLASS}]$ |
| get $\text{REF}.\langle \text{CLASS}, \text{FIELD} \rangle \triangleright \text{CONT}$ | $\text{CONT}[\text{REF} : \text{CLASS.FIELD}]$ |
| set $\text{REF}.\langle \text{CLASS}, \text{FIELD} \rangle := \text{VALUE} \triangleright \text{CONT}$ | $\text{CONT}[\text{REF} : \text{CLASS.FIELD} = \text{VALUE}]$ |
| call $\text{REF}.\text{METHOD}(\overrightarrow{\text{VALUE}}) \triangleright \text{CONT}$ | $\text{CONT}[\text{REF}.\text{METHOD}(\overrightarrow{\text{VALUE}})]$ |
| call $\text{REF}.\langle \text{CLASS}, \text{METHOD} \rangle(\overrightarrow{\text{VALUE}}) \triangleright \text{CONT}$ | super = $\text{REF} : \text{CLASS}.\text{METHOD}(\overrightarrow{\text{VALUE}})$ |
| inspect $\text{REF} \triangleright \text{CONT}[\text{view } [] \leq \text{TNAME REF}]$ | $\text{CONT}[\text{view TNAME REF}]$ |
| error:typecast | $\text{CONT}[\text{view TNAME null}]$ |
| error:typecast | $\text{CONT}[\text{view CLASS} \leq \text{TNAME VALUE}]$ if $\text{CLASS} \not\leq_{PROG}^{CJ} \text{TNAME}$ |
| error:null | $\text{CONT}[\text{null} : \text{CLASS.FIELD}]$ |
| error:null | $\text{CONT}[\text{null} : \text{CLASS.FIELD} = \text{VALUE}]$ |
| error:null | $\text{CONT}[\text{null}.\text{METHOD}(\overrightarrow{\text{VALUE}})]$ |

• $\mapsto_{PROG}^{CJ} : \text{EXPR} \xrightarrow{P} \text{EXPR}$ is:

$$\begin{aligned}
\text{CONT}[\text{let VAR} = \text{VALUE in EXPR}] &\mapsto_{PROG}^{CJ} \text{CONT}[\text{EXPR}[\text{VALUE}/\text{VAR}]] \\
\text{CONT}[\text{view CLASS} \leq \text{TNAME VALUE}] &\mapsto_{PROG}^{CJ} \text{CONT}[\text{VALUE}] \text{ if } \text{CLASS} \leq_{PROG}^{CJ} \text{TNAME}
\end{aligned}$$

Figure 10: Lower level parameters for ClassicJava semantics.

$$\text{METHOD}^{CJ} = \text{METHOD}^\ell + \langle \text{CLASS}^{CJ}, \text{METHOD}^\ell \rangle$$

$$invoke_{PROG}^{CJ}(\text{REF}, \text{OBJECT}, \langle \text{CLASS}, \text{METHOD} \rangle, \overrightarrow{\text{VALUE}}) = evalmethod_{PROG}^{CJ}(\text{REF}, \text{CLASS}, \text{METHOD}, \overrightarrow{\text{VALUE}})$$

$$eval_{PROG}^{CJ}(\text{EXPR}) = \begin{cases} \dots \\ \text{call } \text{REF}.\langle \text{CLASS}, \text{METHOD} \rangle(\overrightarrow{\text{VALUE}}) \triangleright \text{CONT} \\ \text{if } \text{EXPR} \mapsto_{PROG}^* \text{CONT}[\text{super} = \text{REF} : \text{CLASS}.\text{METHOD}(\overrightarrow{\text{VALUE}})] \\ \text{CONT}[evalmethod_{PROG}^{CJ}(\text{REF}, \text{CLASS}, \text{METHOD}, \overrightarrow{\text{VALUE}})] \\ \text{if } \text{EXPR} \mapsto_{PROG}^* \text{CONT}[\text{super} = \text{REF} : \text{CLASS}.\text{METHOD}(\overrightarrow{\text{VALUE}})] \end{cases}$$

Figure 11: Changes for an alternate interpretation of ClassicJava.

ClassicJava syntax. Programs in ClassicJava are a sequence of definitions followed by a method invocation expression; this resembles the call to **main** in Java. Continuations are evaluation contexts over expressions. An object’s static record is the name of its class. Field names include a field label and a class name. Method names include a label and optionally a class name. The sole primitive value is **null**. ClassicJava defines errors for method invocation, null dereference, and failed typecasts.

Our representation of ClassicJava is nearly identical to the original; there are only two changes. We restrict the form of the initial expression to invoke a method of no arguments on a newly constructed object. By making this standalone expression a special case, all meaningful computation takes place within an object.

The other difference lies in the **view** expression. A typecast operation requires an **inspect** message to access an object’s runtime type; the message’s continuation must accept the object’s class and compare it to a static type. ClassicJava has no appropriate evaluation context. We add an elaborated form of typecasts to expressions and continuations. This form adds an explicit position for the class of the object. Our semantics uses this as an intermediate step in evaluating **view** expressions.

Figure 10 defines the semantics of ClassicJava relative to the type relations described in Figure 9. Program initialization installs the receiver of the initial method call as the controlling object and issues the method call. The $resume^{CJ}$ and $invoke^{CJ}$ functions are defined in terms of total functions $evalmethod^{CJ}$, $eval^{CJ}$, and partial function \mapsto^{CJ} .

Method invocation uses $evalmethod^{CJ}$ for dispatch. The $evalmethod^{CJ}$ function looks up the appropriate method in the program’s class definitions. It substitutes the method’s receiver and parameters, then calls $eval^{CJ}$ to evaluate the expression.

The ClassicJava step relation (\mapsto^{CJ}) performs all purely object-internal computations. It reduces **let** expressions by substitution and completes successful typecasts by replacing an elaborated **view** expression with its argument.

The $eval^{CJ}$ function first performs all internal reductions on its argument, then operates by cases on the result. Object creation, field lookup and mutation, method calls and method returns all generate framework actions of the same kind. Unelaborated **view** expressions produce inspection actions, using the elaborated typecast expression as shown above. The $eval^{CJ}$ function signals an error for all null dereferences and typecast failures.

Calls to an object’s superclass generate method call actions; the method name includes the superclass name for method dispatch, which distinguishes it from the current method.

4.2.1 Alternate Interpretation of ClassicJava

Our parameterization of the sequence trace framework for ClassicJava answers the question: “what parts of ClassicJava are object-oriented?” In the semantics above, the answer is clear: object creation, field lookup and mutation, method calls, method returns, and superclass method calls.

We can now consider the merits of this interpretation. The most debatable result is superclass method calls. They take place entirely inside one object and cannot be invoked by outside objects, yet we have formalized them as messages. An alternate perspective might formulate superclass method

calls as object-internal computation for comparison.

Our framework is flexible enough to allow this reinterpretation of ClassicJava. In the prior semantics, $eval^{CJ}$ turns a **super** expression into a method call. A method call uses $invoke^{CJ}$ which uses $evalmethod^{CJ}$. We can rewrite $eval^{CJ}$ to call $evalmethod^{CJ}$ directly in the **super** rule. No object-oriented action is created. The extra clauses for $METHOD^{CJ}$ and $evalmethod^{CJ}$ that were used for superclass calls can be removed. These modifications to ClassicJava are shown in Figure 11.

Now that we have two different semantics for ClassicJava, we can compare them to see what the tradeoffs are; different applications of the ClassicJava model can use either interpretations as appropriate.

4.3 Multiple Dispatch for Sequence Traces

Method dispatch is an important feature of many object-oriented languages. They extend the usual receiver dispatch with some form of *multiple dispatch*, in which a method’s execution is determined by its arguments as well as its receiver. The λ -calculus of Castagna et alii [5] provides a model for single-argument dispatch. We sketch a variant on the λ -calculus in the sequence trace framework. We can then simulate multiple dispatch with higher-order functions or extend the language with multiple function arguments to achieve multiple dispatch.

Figure 4.3 shows the syntax of our dispatch calculus. It extends the λ -calculus with type annotations on functions, an arbitrary set of typed primitive values, and overload expressions which join several typed functions together.

Programs in this setting are expressions. Continuations are evaluation contexts. Objects are functions of both the simple and overloaded kind; an overloaded function’s static record maps each branch’s input type to the branch’s implementation and result type. The language has a single method: function application. The dispatch calculus assumes an arbitrary set of typed primitive values.

Expressions in this language include variables, values, simple and overloaded function abstractions, function application and dispatch. Simple function expressions are annotated with input and output types. Overloaded functions contain a number of function values along with their input and output types. Function application contains a function and its argument. Dispatch expressions are generated during evaluation; they represent the application of an overloaded function to an argument of an explicit type.

We provide only an informal discussion of the operations in this language. The evaluation of function and overload expressions constructs new objects. Function application invokes the operator’s **apply** method on the argument. We split dispatch on overloaded functions into two steps. The **typeof** operator determines the type of an argument, issuing an **inspect** message when given a reference. The dispatch operation itself is given the set of branches, the argument, and the argument’s type; it applies the single most appropriate branch to the argument or signals a dispatch error.

Operations such as multiple dispatch are not inherently object-oriented. We have implemented argument dispatch entirely internal to objects without adding anything to the upper level of our framework. In general, we can interpret new object-oriented features without resorting to changing our framework.

| | | | |
|-------------------------|-----|--|---|
| $\text{PROG}^{\&}$ | $=$ | $\text{EXPR}^{\&}$ | Programs |
| $\text{CONT}^{\&}$ | $=$ | $\begin{aligned} &[] \\ & (\text{CONT}^{\&} \text{EXPR}^{\&}) \\ & (\text{VALUE}^{\&} \text{CONT}^{\&}) \\ & (\text{OVERLOAD}^{\&} (\mathbf{typeof} \text{CONT}^{\&}) \text{VALUE}^{\&}) \\ & \frac{[\text{VALUE}^{\&} : \text{TYPE}^{\&} \longrightarrow \text{TYPE}^{\&} \quad \text{CONT}^{\&} : \text{TYPE}^{\&} \longrightarrow \text{TYPE}^{\&}]}{\text{EXPR}^{\&} : \text{TYPE}^{\&} \longrightarrow \text{TYPE}^{\&}} \end{aligned}$ | Empty Context Function Context Argument Context Type Context Overload Context |
| $\text{STATIC}^{\&}$ | $=$ | $\text{FUNCTION}^{\&} \text{OVERLOAD}^{\&}$ | Static Record |
| $\text{FIELD}^{\&}$ | $=$ | \emptyset | Field Names |
| $\text{METHOD}^{\&}$ | $=$ | \mathbf{apply} | Method Name |
| $\text{PRIM}^{\&}$ | $=$ | typed constants | Primitive Values |
| $\text{LANGERROR}^{\&}$ | $=$ | $\mathbf{error:method} \mathbf{error:dispatch}$ | $\lambda\&$ -calculus Errors |
| where | | | |
| $\text{EXPR}^{\&}$ | $=$ | $\begin{aligned} &\text{VAR}^{\text{TYPE}^{\&}} \\ & \text{VALUE} \\ & \frac{\lambda \text{VAR}^{\text{TYPE}^{\&}} : \text{TYPE}^{\&}. \text{EXPR}^{\&}}{[\text{EXPR}^{\&} : \text{TYPE}^{\&} \longrightarrow \text{TYPE}^{\&}] \text{ (EXPR}^{\&} \text{ EXPR}^{\&})} \\ & (\text{EXPR}^{\&} \text{TEXPR}^{\&} \text{EXPR}^{\&}) \end{aligned}$ | Variable Value Typed Function Overloaded Function Function Application Overload Dispatch |
| $\text{TEXPR}^{\&}$ | $=$ | $\text{TYPE}^{\&} (\mathbf{typeof} \text{EXPR}^{\&})$ | Type Expressions |
| $\text{TYPE}^{\&}$ | $=$ | $\begin{aligned} &\text{BASE}^{\&} \\ & \frac{[\text{TYPE}^{\&} \longrightarrow \text{TYPE}^{\&}]}{\langle \text{VAR}^{\text{TYPE}^{\&}}, \text{TYPE}^{\&}, \text{EXPR}^{\&} \rangle} \end{aligned}$ | Base Type Overloaded Function Type |
| $\text{FUNCTION}^{\&}$ | $=$ | $\langle \text{VAR}^{\text{TYPE}^{\&}}, \text{TYPE}^{\&}, \text{EXPR}^{\&} \rangle$ | Function Record |
| $\text{OVERLOAD}^{\&}$ | $:$ | $\text{TYPE}^{\&} \xrightarrow{f} \langle \text{TYPE}^{\&}, \text{VALUE}^{\&} \rangle$ | Overload Record |

Figure 12: Lower-level parameters for modified $\lambda\&$ -calculus syntax.

5. SEQUENCE TRACE SOUNDNESS

Our two-level semantic framework comes with a two-level type system. The purpose of this type system is to eliminate *all* upper-level type errors (reference error, field error) and to allow only those overall on which the lower-level insist. For example, in the case of ClassicJava, the lower level cannot rule out `null` pointer errors and must therefore raise the relevant exceptions.

Each level defines a portion of the type judgments. The upper level relies on the lower level judgments and vice versa. Naturally, the uniform type soundness theorem holds only if the lower-level elements of the type system satisfy certain constraints.

The lower level parameters include several sets, functions, and type judgments, shown together with the constraints they must satisfy in Figure 13. The set `TYPE` defines types for the languages values; `OBJTYPE` defines the subset of `TYPE` representing the types of objects. The subset `OK-ERROR` of `LANGERROR` distinguishes the runtime exceptions that well-typed programs may throw.

The relation \leq induces a partial order on types. The total functions *fields* and *methods* define the field and method signatures of object types. The total function *metatype* determines the type of a static record from the type of its container object; it is needed to type **inspect** messages.

The lower level type system must provide type judgments for programs, continuations, and primitive values. It must also provide a type judgment for static records when used in the context of an object record (as opposed to their type as a value).

The type rules `INIT`, `RESUME`, and `INVOKE` constrain the lower-level framework functions of the same names. In addition, a sound system requires all three to be total functions, whereas the untyped operational semantics allows *resume* and *invoke* to be partial. The lower level type system must guarantee these rules, while the upper level relies on them for a parametric soundness proof.

The upper-level type system defines type judgments for everything else: program states, object pools, stacks, references, static records (as values), object records, dynamic records, and actions of both the message and error variety. Figure 14 shows the full set of upper-level judgments and type rules.

5.1 Framework Soundness

The type system satisfies a conventional type soundness theorem. Its statement, however, assumes that lower-level exceptions are typed.

THEOREM 1. *If $\vdash^{\ell} \text{PROG} : \text{TYPE}$, then either $\text{PROG} \uparrow$ (diverges) or $\text{PROG} \downarrow \text{RESULT}$ and $\text{PROG} \vdash^u \text{RESULT} : \text{TYPE}$.*

PROOF. The proof uses progress and preservation lemmas, adapting the method of Wright and Felleisen [36]. By the `INIT` rule, the initial state of any program preserves the program's type. Lemmas 1 and 2 establish the soundness of evaluation at all subsequent states. \square

LEMMA 1. *If $\text{PROG} \vdash^u \text{STATE} : \text{TYPE}$, then either STATE is a result state or $\text{STATE} \mapsto_{\text{PROG}} \text{STATE}'$ for some STATE' .*

PROOF. The proof follows by cases on `STATE`. The step relation \mapsto has a clause for every non-result state. \square

Sets and functions:

| | | |
|--------------------------|--|---|
| TYPE | any set | Value types |
| OBJTYPE | \subseteq TYPE | Object types |
| OKERROR | \subseteq LANGERROR | Allowable runtime errors |
| \leq_{PROG} | partial order on TYPE | Subtype relation |
| $fields_{\text{PROG}}$ | $: \text{OBJTYPE} \longrightarrow (\text{FIELD} \xrightarrow{f} \text{TYPE})$ | Produces an object type's field signatures. |
| $methods_{\text{PROG}}$ | $: \text{OBJTYPE} \longrightarrow (\text{METHOD} \xrightarrow{f} \langle \overline{\text{TYPE}}, \text{TYPE} \rangle)$ | Produces an object type's method signatures. |
| $metatype_{\text{PROG}}$ | $: \text{OBJTYPE} \longrightarrow \text{TYPE}$ | Given an object's type, produces the type of its static record. |

Type judgments:

| | |
|--|---|
| $\vdash^\ell \text{PROG} : \text{TYPE}$ | Program PROG has type TYPE. |
| $\text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_1 \xrightarrow{c} \text{TYPE}_2$ | Continuation CONT produces an action of type TYPE ₂ when given input of type TYPE ₁ . |
| $\text{PROG}, \text{POOL} \vdash_0^\ell \text{STATIC} : \text{OBJTYPE}$ | Static record STATIC is well-formed in an object of type OBJTYPE. |
| $\text{PROG}, \text{POOL} \vdash_v^\ell \text{PRIM} : \text{TYPE}$ | Primitive value PRIM has type TYPE. |

Satisfying:

$$\begin{array}{c}
 \text{INIT} \frac{\vdash^\ell \text{PROG} : \text{TYPE}}{\text{PROG} \vdash^u \text{init}(\text{PROG}) : \text{TYPE}} \\
 \\
 \text{RESUME} \frac{\text{PROG}, \text{POOL} \vdash_v^{u\ell} \text{VALUE} : \text{TYPE}_1 \quad \text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_2 \xrightarrow{c} \text{TYPE}_3 \quad \text{TYPE}_1 \leq_{\text{PROG}} \text{TYPE}_2 \quad \text{TYPE}_4 \leq_{\text{PROG}} \text{TYPE}_3}{\text{PROG}, \text{POOL} \vdash_a^u \text{resume}_{\text{PROG}}(\text{CONT}, \text{VALUE}) : \text{TYPE}_4} \\
 \\
 \text{INVOKE} \frac{\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE} \quad \text{PROG}, \text{POOL} \vdash_v^{u\ell} \overline{\text{VALUE}} : \overline{\text{TYPE}}' \quad \text{methods}_{\text{PROG}}(\text{OBJTYPE})(\text{METHOD}) = \langle \overline{\text{TYPE}}'', \text{TYPE}_1 \rangle \quad \text{TYPE}' \leq_{\text{PROG}} \text{TYPE}'' \quad \text{TYPE}_2 \leq_{\text{PROG}} \text{TYPE}_1}{\text{PROG}, \text{POOL} \vdash_a^u \text{invoke}(\text{REF}, \text{POOL}(\text{REF}), \text{METHOD}, \overline{\text{VALUE}}) : \text{TYPE}_2}
 \end{array}$$

Figure 13: Lower level parameters to the type system.

Judgments:

| | |
|--|---|
| $\text{PROG} \vdash^u \text{STATE} : \text{TYPE}$ | State STATE has type TYPE. |
| $\text{PROG} \vdash^u \text{POOL}$ | Object pool POOL is well-formed. |
| $\text{PROG}, \text{POOL} \vdash^u \text{STACK} : \text{TYPE}_1 \xrightarrow{s} \text{TYPE}_2$ | Stack STACK represents computation producing type TYPE_2 if the current method produces type TYPE_1 . |
| $\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE}$ | Reference REF has type OBJTYPE. |
| $\text{PROG}, \text{POOL} \vdash_v^u \text{STATIC} : \text{TYPE}$ | Static record STATIC has type TYPE as a value. |
| $\text{PROG}, \text{POOL} \vdash_o^u \text{OBJECT} : \text{OBJTYPE}$ | Object record OBJECT is well-formed in an object of type OBJTYPE. |
| $\text{PROG}, \text{POOL} \vdash_o^u \text{DYNAMIC} : \text{OBJTYPE}$ | Dynamic record DYNAMIC is well-formed in an object of type OBJTYPE. |
| $\text{PROG}, \text{POOL} \vdash_a^u \text{ACTION} : \text{TYPE}$ | Action ACTION represents computation in a method producing type TYPE. |

Type rules:

| | | | |
|------------|--|---|---|
| | $\text{PROG} \vdash^u \text{POOL}$ | $\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE}$ | |
| STATE | $\frac{\text{PROG}, \text{POOL} \vdash_a^u \text{ACTION} : \text{TYPE}_1 \quad \text{PROG}, \text{POOL} \vdash^u \text{STACK} : \text{TYPE}_2 \xrightarrow{s} \text{TYPE}_3 \quad \text{TYPE}_1 \leq_{\text{PROG}} \text{TYPE}_2}{\text{PROG} \vdash^u \langle \text{POOL}, \text{STACK}, \text{REF}, \text{ACTION} \rangle : \text{TYPE}_3}$ | | |
| POOL | $\frac{\forall \text{REF} \in \text{dom}(\text{POOL}). \text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE}}{\text{PROG} \vdash^u \text{POOL}}$ | EMPTYSTACK | $\frac{}{\text{PROG}, \text{POOL} \vdash^u \epsilon : \text{TYPE} \xrightarrow{s} \text{TYPE}}$ |
| STACKFRAME | $\frac{\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE} \quad \text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_1 \xrightarrow{c} \text{TYPE}_2 \quad \text{PROG}, \text{POOL} \vdash^u \text{STACK} : \text{TYPE}_3 \xrightarrow{s} \text{TYPE}_4 \quad \text{TYPE}_2 \leq_{\text{PROG}} \text{TYPE}_3}{\text{PROG}, \text{POOL} \vdash^u \langle \text{REF}, \text{CONT} \rangle \text{STACK} : \text{TYPE}_1 \xrightarrow{s} \text{TYPE}_4}$ | | |
| METATYPE | $\frac{\text{PROG}, \text{POOL} \vdash_o^\ell \text{STATIC} : \text{OBJTYPE}}{\text{PROG}, \text{POOL} \vdash_v^u \text{STATIC} : \text{metatype}_{\text{PROG}}(\text{OBJTYPE})}$ | REFERENCE | $\frac{\text{POOL}(\text{REF}) = \text{OBJECT} \quad \text{PROG}, \text{POOL} \vdash_o^u \text{OBJECT} : \text{OBJTYPE}}{\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE}}$ |
| OBJECT | $\frac{\text{PROG}, \text{POOL} \vdash_o^\ell \text{STATIC} : \text{OBJTYPE} \quad \text{PROG}, \text{POOL} \vdash_o^u \text{DYNAMIC} : \text{OBJTYPE}}{\text{PROG}, \text{POOL} \vdash_o^u \text{STATIC DYNAMIC} : \text{OBJTYPE}}$ | | |
| DYNAMIC | $\frac{\text{dom}(\text{DYNAMIC}) = \text{dom}(\text{fields}_{\text{PROG}}(\text{OBJTYPE})) \quad \forall \text{FIELD} \in \text{dom}(\text{DYNAMIC}). \text{PROG}, \text{POOL} \vdash_v^{u\ell} \text{VALUE} : \text{TYPE} \wedge \text{TYPE} \leq_{\text{PROG}} \text{fields}_{\text{PROG}}(\text{OBJTYPE})(\text{FIELD})}{\text{PROG}, \text{POOL} \vdash_o^u \text{DYNAMIC} : \text{OBJTYPE}}$ | | |
| NEW | $\frac{\text{PROG}, \text{POOL} \vdash_o^u \text{OBJECT} : \text{OBJTYPE} \quad \text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_1 \xrightarrow{c} \text{TYPE}_2 \quad \text{OBJTYPE} \leq_{\text{PROG}} \text{TYPE}_1}{\text{PROG}, \text{POOL} \vdash_a^u \text{new OBJECT} \triangleright \text{CONT} : \text{TYPE}_2}$ | INSPECT | $\frac{\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE} \quad \text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_1 \xrightarrow{c} \text{TYPE}_2 \quad \text{metatype}_{\text{PROG}}(\text{OBJTYPE}) \leq_{\text{PROG}} \text{TYPE}_1}{\text{PROG}, \text{POOL} \vdash_a^u \text{inspect REF} \triangleright \text{CONT} : \text{TYPE}_2}$ |
| GET | $\frac{\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE} \quad \text{fields}_{\text{PROG}}(\text{OBJTYPE})(\text{FIELD}) = \text{TYPE}_1 \quad \text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_2 \xrightarrow{c} \text{TYPE}_3 \quad \text{TYPE}_1 \leq_{\text{PROG}} \text{TYPE}_2}{\text{PROG}, \text{POOL} \vdash_a^u \text{get REF.FIELD} \triangleright \text{CONT} : \text{TYPE}_3}$ | SET | $\frac{\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE} \quad \text{PROG}, \text{POOL} \vdash_v^{u\ell} \text{VALUE} : \text{TYPE}_1 \quad \text{fields}_{\text{PROG}}(\text{OBJTYPE})(\text{FIELD}) = \text{TYPE}_2 \quad \text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_3 \xrightarrow{c} \text{TYPE}_4 \quad \text{TYPE}_1 \leq_{\text{PROG}} \text{TYPE}_2 \quad \text{TYPE}_2 \leq_{\text{PROG}} \text{TYPE}_3}{\text{PROG}, \text{POOL} \vdash_a^u \text{set REF.FIELD} := \text{VALUE} \triangleright \text{CONT} : \text{TYPE}_4}$ |
| CALL | $\frac{\text{PROG}, \text{POOL} \vdash_v^u \text{REF} : \text{OBJTYPE} \quad \text{PROG}, \text{POOL} \vdash_v^{u\ell} \text{VALUE} : \text{TYPE}' \xrightarrow{\quad} \text{methods}_{\text{PROG}}(\text{OBJTYPE})(\text{METHOD}) = \langle \text{TYPE}'', \text{TYPE}_1 \rangle \quad \text{PROG}, \text{POOL} \vdash^\ell \text{CONT} : \text{TYPE}_2 \xrightarrow{c} \text{TYPE}_3 \quad \text{TYPE}' \leq_{\text{PROG}} \text{TYPE}'' \quad \text{TYPE}_1 \leq_{\text{PROG}} \text{TYPE}_2}{\text{PROG}, \text{POOL} \vdash_a^u \text{call REF.METHOD}(\text{VALUE}) \triangleright \text{CONT} : \text{TYPE}_3}$ | | |
| RETURN | $\frac{\text{PROG}, \text{POOL} \vdash_v^{u\ell} \text{VALUE} : \text{TYPE}}{\text{PROG}, \text{POOL} \vdash_a^u \text{return VALUE} : \text{TYPE}}$ | ERROR | $\frac{}{\text{PROG}, \text{POOL} \vdash_a^u \text{OKERROR} : \text{TYPE}}$ |

Figure 14: Upper level definition of the type system.

LEMMA 2. *If $\text{PROG} \vdash^u \text{STATE} : \text{TYPE}$ and $\text{STATE} \mapsto_{\text{PROG}} \text{STATE}'$, then there exists TYPE' such that $\text{PROG} \vdash^u \text{STATE}' : \text{TYPE}'$ and $\text{TYPE}' \leq_{\text{PROG}} \text{TYPE}$.*

PROOF. The proof proceeds by cases on the clauses in the definition of \mapsto . See Figure 4 for reference. These cases use lemmas 3 and 4, which establish type preservation under modified object pools.

Object Creation. The type of REF' in the extended pool is the same as the type of OBJECT in POOL , by the REFERENCE rule and lemma 3. The type derivation of the **new** message proves that CONT accepts a supertype of this as input, so by the RESUME rule the action's type is preserved. Lemma 3 proves that the rest of the state's type is preserved.

Object Inspection. From the action's type derivation we know that REF has type OBJTYPE and CONT 's input is a supertype of $\text{metatype}_{\text{PROG}}(\text{OBJTYPE})$. We know $\text{PROG}, \text{POOL} \vdash_{\circ}^{\ell} \text{STATIC} : \text{OBJTYPE}$ from the REFERENCE and OBJECT rules; therefore METATYPE gives us $\text{PROG}, \text{POOL} \vdash_{\circ}^u \text{STATIC} : \text{metatype}_{\text{PROG}}(\text{OBJTYPE})$. We can apply the RESUME rule to see that the action's type is preserved; the rest of the state is unchanged.

Field Lookup. From the action's type derivation, we know REF has type OBJTYPE and the signature for FIELD in OBJTYPE is a subtype of CONT 's input. By the type derivation of REF and the DYNAMIC rule, we know that VALUE is a suitable input for CONT . The RESUME rule proves the action's type is preserved; the rest of the state is unchanged.

Field Mutation. The **set** action's type derivation proves that REF has type OBJTYPE , that VALUE provides an appropriate type for FIELD in OBJTYPE , and that CONT accepts input of that type. By lemma 4, the modified object pool is well-formed and the types of the stack, object reference, and continuation are preserved. The type of the action is preserved due to the RESUME rule.

Method Call. The type derivation of STATE proves that REF has type OBJTYPE , STACK accepts CONT 's output, that CONT accepts the output of METHOD in OBJTYPE , and that the method in turn accepts the argument types given to it. The STACKFRAME rule proves the new stack is well-typed and INVOKE proves the new action is appropriate for the stack. The overall state's type is preserved.

Method Return. From the first state's type derivation, we know that STACK accepts CONT 's output and that CONT 's input type accepts VALUE . By the RESUME rule, the new stack accepts the type of the new action. The STATE rule establishes that the state's type is preserved.

Reference Error. The type rules INSPECT , GET , SET , and CALL all require the receiver reference be typed; the REFERENCE rule dictates that typed references have object pool entries. Therefore the type system rejects reference errors statically. This case is impossible in well-typed programs.

Field Error. The GET and SET rules require that FIELD has a signature in the receiver's object type. By the rules REFERENCE , OBJECT , and DYNAMIC the field must therefore exist in the dynamic record. All field errors are rejected statically. This case is impossible in a well-typed program.

With all the cases proved, the type of STATE is preserved in the transition to STATE' . \square

LEMMA 3. *If*

$$\begin{aligned} & \text{PROG} \vdash^u \text{POOL}, \\ & \text{REF} \notin \text{POOL}, \text{ and} \\ & \text{PROG}, \text{POOL} \vdash_{\circ}^u \text{OBJECT} : \text{OBJTYPE}, \end{aligned}$$

then $\text{PROG} \vdash^u \text{POOL}[\text{REF} \mapsto \text{OBJECT}]$. Additionally, all type derivations valid with respect to the first pool remain valid with respect to the new pool.

PROOF. The proof of well-formedness of the new pool follows from the POOL , REFERENCE , and OBJECT rules. The proof of preserved types relies on the fact that the object pool is only used directly in the REFERENCE and OBJECT rules. Types derived from those rules are preserved; therefore so are all others. \square

LEMMA 4. *If*

$$\begin{aligned} & \text{PROG} \vdash^u \text{POOL}, \\ & \text{PROG}, \text{POOL} \vdash_{\circ}^u \text{REF} : \text{OBJTYPE}, \\ & \text{fields}_{\text{PROG}}(\text{OBJTYPE})(\text{FIELD}) = \text{TYPE}_1, \\ & \text{PROG}, \text{POOL} \vdash_{\circ}^{\ell} \text{VALUE} : \text{TYPE}_2, \\ & \text{TYPE}_2 \leq_{\text{PROG}} \text{TYPE}_1, \text{ and} \\ & \text{POOL}(\text{REF}) = \text{STATIC DYNAMIC}, \end{aligned}$$

then

$$\text{PROG} \vdash^u \text{POOL}[\text{REF} \mapsto \text{STATIC DYNAMIC}[\text{REF} \mapsto \text{VALUE}]].$$

Additionally, all type derivations valid with respect to the first pool remain valid with respect to the new pool.

PROOF. The proof is a slight variation on lemma 3. It uses the DYNAMIC and OBJECT rules to prove that the new object record is well-typed and the REFERENCE rule to prove that references to the object are still well-typed. \square

5.2 ClassicJava Soundness

Figure 15 presents the lower-level definitions of the ClassicJava type system. To obtain a type soundness result, we must prove that these definitions satisfy the criteria for soundness; together with the upper-level soundness theorem, this establishes a result analogous to Flatt et alii's ClassicJava soundness proof [13].

Types in this formulation of ClassicJava are class or interface names or a **meta**-type, representing the static record for an object of a given type. Object types must be class or interface names. ClassicJava allows null dereferences and failed typecasts in well-typed programs. The subtyping relation is unchanged from the original definition; **meta**-types have no relationship with other types.

The $\text{fields}^{\text{CJ}}$ and $\text{methods}^{\text{CJ}}$ functions are defined in terms of the \in^{CJ} relation. The $\text{metatype}^{\text{CJ}}$ function defines the static record of CLASS to have type **meta CLASS**.

The type rules are defined in terms of the original ClassicJava judgments for typing programs and expressions. We

| | | | |
|----------------------------------|-----|--|----------------|
| TYPE^{CJ} | $=$ | $\text{TNAME}^{\text{CJ}} \mid \mathbf{meta} \text{TNAME}^{\text{CJ}}$ | Types |
| $\text{OBJTYPE}^{\text{CJ}}$ | $=$ | TNAME^{CJ} | Object Types |
| $\text{OKERROR}^{\text{CJ}}$ | $=$ | $\mathbf{error:null} \mid \mathbf{error:typecast}$ | Runtime Errors |
| $\leq_{\text{PROG}}^{\text{CJ}}$ | | as used in Section 4.2 | Subtyping |

$$\begin{aligned}
\mathit{fields}_{\text{PROG}}^{\text{CJ}}(\text{TNAME}) &= [(\text{CLASS}, \text{FIELD}) \mapsto \text{TYPE} \mid \langle \text{CLASS}, \text{FIELD}, \text{TYPE} \rangle \in_{\text{PROG}}^{\text{CJ}} \text{TNAME}] \\
\mathit{methods}_{\text{PROG}}^{\text{CJ}}(\text{TNAME}) &= [(\text{CLASS}, \text{METHOD}) \mapsto \langle \overline{\text{TYPE}}, \overline{\text{TYPE}} \rangle \mid \text{TNAME} \leq_{\text{PROG}}^{\text{CJ}} \text{CLASS} \\
&\quad \text{and } \langle \text{METHOD}, \overline{\text{TYPE}} \rangle \longrightarrow \text{TYPE}, \overline{\text{VAR}}, \overline{\text{EXPR}} \in_{\text{PROG}}^{\text{CJ}} \text{CLASS}] \\
&\quad [\text{METHOD} \mapsto \langle \overline{\text{TYPE}}, \overline{\text{TYPE}} \rangle \mid \langle \text{METHOD}, \overline{\text{TYPE}} \rangle \longrightarrow \text{TYPE}, \overline{\text{VAR}}, \overline{\text{EXPR}} \in_{\text{PROG}}^{\text{CJ}} \text{TNAME}] \\
\mathit{metatype}_{\text{PROG}}^{\text{CJ}}(\text{TNAME}) &= \mathbf{meta} \text{TNAME}
\end{aligned}$$

$$\text{PROG} \frac{\vdash_{\text{p}}^{\text{CJ}} \text{PROG} : \text{OBJTYPE}}{\vdash^{\ell} \text{PROG} : \text{OBJTYPE}} \quad \text{CONT} \frac{\text{VAR} \notin \text{CONT} \quad \text{PROG}, \mathit{getenv}^{\text{CJ}}(\text{POOL})[\text{VAR} \mapsto \text{TYPE}_1] \vdash_{\text{e}}^{\text{CJ}} \text{CONT}[\text{VAR}] : \text{TYPE}_2}{\text{PROG}, \text{POOL} \vdash^{\ell} \text{CONT} : \text{TYPE}_1 \xrightarrow{c} \text{TYPE}_2}$$

$$\text{CLASS} \frac{}{\text{PROG}, \text{POOL} \vdash_{\circ}^{\ell} \text{CLASS} : \text{CLASS}}$$

$$\text{NULL} \frac{}{\text{PROG}, \text{POOL} \vdash_{\vee}^{\ell} \mathbf{null} : \text{OBJTYPE}}$$

Original ClassicJava judgments:

| | |
|---|--|
| ENV | ClassicJava environment. |
| $\vdash_{\text{p}}^{\text{CJ}} \text{PROG} : \text{TYPE}$ | ClassicJava program PROG has type TYPE. |
| $\text{PROG}, \text{ENV} \vdash_{\text{e}}^{\text{CJ}} \text{EXPR} : \text{TYPE}$ | ClassicJava expression EXPR has type TYPE. |

where

$$\text{VIEWSUB}^{\text{CJ}} \frac{\text{PROG}, \text{ENV} \vdash_{\text{e}}^{\text{CJ}} \text{EXPR} : \mathbf{meta} \text{TNAME}' \quad \text{PROG}, \text{ENV} \vdash_{\text{e}}^{\text{CJ}} \text{VALUE} : \text{TNAME}'}{\text{PROG}, \text{ENV} \vdash_{\text{e}}^{\text{CJ}} \mathbf{view} \text{EXPR} \leq \text{TNAME} \text{ VALUE} : \text{TNAME}}$$

and $\mathit{getenv}^{\text{CJ}} : \text{Pool} \longrightarrow \text{Env}$ is:

$$\mathit{getenv}^{\text{CJ}}(\text{POOL}) = [\text{REF} \mapsto \text{CLASS} \mid \text{POOL}(\text{REF}) = \text{CLASS DYNAMIC}]$$

Figure 15: Lower level parameters for the ClassicJava type system.

extend the ClassicJava rules to cover the elaborated form of **view** and define a function $getenv^{CJ}$ to construct a ClassicJava environment from the object pool.

THEOREM 2. *The sequence trace representation of ClassicJava is type-sound.*

PROOF. The goal of our proof is to validate that all the lower-level type relations and functions satisfy the three constraints from the preceding section. We dedicate one lemma to each constraint. Specifically, lemma 10 concerns $init^{CJ}$, lemma 12 is about $invoke^{CJ}$, and lemma 11 establishes $resume^{CJ}$. \square

Before we can prove anything meaningful about ClassicJava evaluation as presented in Figure 10, we must prove that the primitive operations of substitution for a variable, plugging expressions into evaluation contexts, \mapsto^{CJ} , $eval^{CJ}$, and $evalmethod^{CJ}$ all preserve the types of their inputs.

LEMMA 5. *If*

$$\begin{aligned} & \text{PROG, ENV}[\text{VAR} \mapsto \text{TYPE}_0] \vdash_{\underline{e}}^{CJ} \text{EXPR} : \text{TYPE}_1, \\ & \text{VAR} \notin \text{dom}(\text{ENV}), \text{ and} \\ & \text{PROG, ENV} \vdash_{\underline{e}}^{CJ} \text{VALUE} : \text{TYPE}_0, \end{aligned}$$

then $\text{PROG, ENV} \vdash_{\underline{e}}^{CJ} \text{EXPR}[\text{VALUE}/\text{VAR}] : \text{TYPE}_2$ *such that* $\text{TYPE}_2 \leq_{\text{PROG}} \text{TYPE}_1$.

PROOF. This proof proceeds by induction on the structure of the type derivation for EXPR. \square

LEMMA 6. *If*

$$\begin{aligned} & \text{PROG, POOL} \vdash_{\underline{o}}^{\ell} \text{CONT} : \text{TYPE}_2 \xrightarrow{c} \text{TYPE}_3, \\ & \text{PROG, } \overline{getenv}^{CJ}(\text{POOL}) \vdash_{\underline{e}}^{CJ} \text{EXPR} : \text{TYPE}_1, \text{ and} \\ & \text{TYPE}_1 \leq_{\text{PROG}} \text{TYPE}_2, \end{aligned}$$

then $\text{PROG, } \overline{getenv}^{CJ}(\text{POOL}) \vdash_{\underline{e}}^{CJ} \text{CONT}[\text{EXPR}] : \text{TYPE}_4$ *such that* $\text{TYPE}_4 \leq_{\text{PROG}} \text{TYPE}_3$.

PROOF. This is a corollary of lemma 5; because the CONT rule types continuations by plugging in a variable, plugging in an arbitrary expression to the same type derivation is a simple matter of substitution. It is worth noting that none of the evaluation contexts capture free variables, so that is not an issue. \square

LEMMA 7. *If*

$$\begin{aligned} & \text{PROG, ENV} \vdash_{\underline{e}}^{CJ} \text{EXPR} : \text{TYPE} \text{ and} \\ & \text{EXPR} \mapsto_{\text{PROG}}^{CJ} \overline{\text{EXPR}}', \end{aligned}$$

then $\text{PROG, ENV} \vdash_{\underline{e}}^{CJ} \overline{\text{EXPR}}' : \text{TYPE}'$ *such that* $\text{TYPE}' \leq_{\text{PROG}} \text{TYPE}$.

PROOF. There are two cases to examine:

Let Binding. This case follows directly from lemma 5.

Successful Typecast. By the type derivation of an elaborated typecast, VALUE has type TNAME' , and the expression is replaced with VALUE only when

$$\text{TNAME}' \leq_{\text{PROG}}^{CJ} \text{TNAME}.$$

Therefore the type TNAME of the expression is preserved.

\square

LEMMA 8. *If*

$$\text{PROG, } \overline{getenv}^{CJ}(\text{POOL}) \vdash_{\underline{e}}^{CJ} \text{EXPR} : \text{TYPE}$$

then

$$\text{PROG, POOL} \vdash_a^u \overline{eval}_{\text{PROG}}^{CJ}(\text{EXPR}) : \text{TYPE}'$$

such that $\text{TYPE}' \leq_{\text{PROG}} \text{TYPE}$.

PROOF. The \overline{eval}^{CJ} function first reduces its input before generating a message. After reduction, \overline{eval}^{CJ} creates an action from the resulting expression. We first argue that the creation of these actions is a total function, then that the function preserves its input's type.

The step function \mapsto^{CJ} strictly reduces the total number of **let** and **view** subexpressions at each step, so there cannot be an infinite sequence of reductions. The clauses of the step relation and the final action generation are exhaustive, so every reduced expression will result in an action. Therefore \overline{eval}^{CJ} is a total function.

The proof that \overline{eval}^{CJ} preserves types proceeds by cases on the fully reduced expression. In each case, the type derivation for the resulting action can be assembled from pieces of the type derivation of the expression; the pieces of the expression are shuffled around but nothing new is generated (there is no substitution or replication). \square

LEMMA 9. *If*

$$\begin{aligned} & \text{PROG, POOL} \vdash_{\underline{o}}^{\ell} \text{STATIC} : \text{TNAME}, \\ & \text{PROG, POOL} \vdash_{\underline{v}}^u \text{REF} : \text{TNAME}, \\ & \text{PROG, POOL} \vdash_{\underline{v}}^{u\ell} \overline{\text{VALUE}} : \text{TYPE}', \\ & \overline{methods}^{CJ}(\text{TNAME})(\text{METHOD}) = (\overline{\text{TYPE}}'', \text{TYPE}_1), \text{ and} \\ & \overline{\text{TYPE}}' \leq_{\text{PROG}} \overline{\text{TYPE}}'', \end{aligned}$$

then $\overline{evalmethod}_{\text{PROG}}^{CJ}(\text{REF}, \text{STATIC}, \text{METHOD}, \overline{\text{VALUE}})$ *has the type* TYPE_2 *such that* $\text{TYPE}_2 \leq_{\text{PROG}} \text{TYPE}_1$.

PROOF. This lemma relies on the type derivation of the program's class definitions. Specifically, it requires that if:

$$\begin{aligned} & (\overline{\text{METHOD}}, \overline{\text{TYPE}}' \longrightarrow \text{TYPE}, \overline{\text{VAR}}, \text{EXPR}) \in_{\text{PROG}}^{CJ} \text{CLASS}, \\ & \text{PROG, ENV} \vdash_{\underline{e}}^{CJ} \overline{\text{VALUE}} : \text{TYPE}'', \\ & \text{PROG, ENV} \vdash_{\underline{e}}^{CJ} \text{REF} : \text{CLASS}', \\ & \overline{\text{TYPE}}'' \leq_{\text{PROG}}^{CJ} \overline{\text{TYPE}}' \text{ and } \text{CLASS}' \leq_{\text{PROG}}^{CJ} \text{CLASS}, \end{aligned}$$

then $\text{PROG, ENV} \vdash_{\underline{e}}^{CJ} \text{EXPR}[\text{REF}/\text{this}][\overline{\text{VALUE}}/\text{VAR}] : \text{TYPE}$. Put simply, when a method body's parameters of appropriate types are filled in, the method must return the right type.

When these conditions hold, $\overline{evalmethod}^{CJ}$ passes the substituted method body to \overline{eval}^{CJ} . By lemma 8, the expression's type is preserved. In a well-typed program, these conditions will always hold; $\overline{evalmethod}^{CJ}$ preserves types and method errors are rejected statically. \square

LEMMA 10. *The INIT rule holds for ClassicJava.*

PROOF. The type of the program

$$\overline{\text{DEF}} \text{ new CLASS.METHOD}()$$

is the result type of METHOD in CLASS. By inspection, we can determine that the initial state's type is the result type of $\overline{methods}_{\text{PROG}}^{CJ}(\text{CLASS})(\text{METHOD})$. By the definition of $\overline{methods}^{CJ}$, these are the same types; therefore \overline{init}^{CJ} preserves the program's type. \square

LEMMA 11. *The RESUME rule holds for ClassicJava.*

PROOF. This is a straightforward application of lemmas 6 and 8. \square

LEMMA 12. *The INVOKE rule holds for ClassicJava.*

PROOF. There are two clauses in $invoke^{CJ}$'s definition; one for simple method names and one for method names specifying a class for dispatch. Both invoke $evalmethod^{CJ}$ on one of REF's (super)classes (as constrained by $methods^{CJ}$); the former clause merely defaults to the most specific class. It suffices to consider the second case.

The premises of the INVOKE rule (plus an application of the REFERENCE and OBJECT rules to determine STATIC's type) allow us to apply lemma 9. \square

6. RELATED WORK

Our work has two inspirational sources. Calculi for communicating processes often model just those actions that relate to process creation, communication, and so on. This corresponds to our isolation of object-oriented actions in the upper level of the framework; of course, our framework also specifies a precise interface for the lower level and, with the specification of a lower level, models entire languages. In the exploration of process calculi, traces play a natural role, though research there often concerns observational equivalence or problems of parallel, concurrent, and even distributed computations. Starting from this insight, Graunke, the second author, and others [15, 16, 24] have recently created a trace calculus for a sequential client-server setting. This calculus models a web client (browser) and web server with the goal of understanding systemic flaws in interactive web programs. Roughly speaking, our paper generalizes Graunke et al.'s prior research to an arbitrarily large and growing pool of objects with a rich set of actions and a wide interface to the underlying computational language.

Others tools for inspecting and debugging program traces exist, tackling the problem from many different perspectives. Lewis [25] presents a so-called omniscient debugger which records every change in program state and reconstructs the execution after the fact. Intermediate steps in the program's execution can thus be debugged even after program completion. This approach is similar to our own, but with emphasis on the pragmatics of debugging rather than presenting an intuitive view of computation.

The work of Walker, Murphy, Steinbok, and Robillard [34] develops a method for constructing abstract interpretations of a trace interactively. They allow a user to write rules grouping primitive program elements into abstract elements. The program trace is reinterpreted with respect to the abstract elements; they provide visualization and summarization tools for the resulting abstract trace.

Richner and Ducasse [32] demonstrate a tool for recovering class collaborations from the execution of object-oriented programs. The process is iterative and interactive; a programmer writes and refines queries to find the meaningful collaborations in a program trace. Their tool presents a browser for collaborations and interactions discovered in the program.

Ducasse, Girba, and Wuyts [10] create regression tests by examining execution traces of legacy software. Queries are written as logic programming predicates; queries that the system satisfies can be kept as tests to apply to updated

versions of the software. This allows the creation of test suites without the need for documentation, understanding the code, or bringing the program into a known state for each test.

Even though our work does not attempt to assign semantics to UML's sequence charts, many pieces of research in this direction exist and share some similarities with our own work. We therefore describe the most relevant pieces of work in this direction here. Lund and Stølen [27] implement an operational semantics for UML 2.0 sequence diagrams with an execution system and a projection system. The execution system handles the representation of the communication medium and the projection system handles choices among lifelines when executing events. They also formalize a type system that separates valid and invalid traces and employs strategies for guiding execution.

Xia and Kane [37] develop a semantics for UML class and sequence diagrams based on attribute grammars. They formalize static attributes that express the objects and operations referenced from within an operation and dynamic attributes that express the relative ordering of execution of operations. They demonstrate how to prove theorems about class and sequence diagrams within their semantics.

Whittle [35] implements a method for generating statecharts from collections of sequence diagrams. Ambiguities and inconsistencies among sequence diagrams are resolved with Object Constraint Language annotations. He lacks a formal semantics and says that his approach is pragmatic.

Okazaki, Aoki, and Katayama [29] also use state machines to combine sequence diagrams. They formulate a new language for state machines called Concurrent Regular Expressions. They present the translation of sequence diagrams into this language and the method for combining the resulting machines. They also present two examples of checking consistency properties of the diagrams using the new state machines.

Li, Liu, and He [26] develop a formal static and dynamic semantics for UML sequence diagrams. Their static semantics relates a sequence diagram to a hierarchical tree structure of method calls and checks consistency against class diagrams. Their dynamic semantics relates sequence diagrams to consistency predicates on the execution of single operations. Much like our sequence trace semantics, their dynamic semantics presents executions as the composition of message passing and small connecting programs within each object.

Nanta,jeewarawat and Sombatsrisomboon [28] define a new model theoretic framework for class and sequence diagrams. They provide conditions for valid interpretations and a set of inference rules for reasoning about diagrams. They demonstrate the use of these rules to infer a class diagram from a sequence diagram.

Hausmann, Heckel, and Sauer [19] introduce a technique dubbed dynamic meta modeling and demonstrate it with a dynamic semantics for executing sequence diagrams directly. Roughly speaking dynamic meta modeling is a collection of graphical rewriting rules on collaboration diagrams.

Cho, Kim, Cha, and Bae [6] present a semantics for sequence diagrams in a new temporal logic called HDTL. They provide a method for computing HDTL formulae from a given sequence diagram and a method for validating the formulae. The validation produces instances of the sequence diagram from traces of many independently operating ob-

jects.

Bernardi, Donatelli, and Merseguer [3] provide a translation from statecharts and sequence diagrams to Labeled Generalized Stochastic Petri Nets. The partial order on messages expressed by a sequence diagram is preserved, as are the synchronous/asynchronous and delayed/instantaneous properties of each action. The resulting Petri nets are used to compute logical properties of the diagrams as well as predict performance results.

7. CONCLUSIONS AND FUTURE WORK

The paper presents a two-level semantics framework for object-oriented programming. The framework carefully distinguishes actions on objects from internal computations of objects. The two levels are separated via a collection of mathematical entities: an unspecified set plus a number of functions. At this point the framework can easily handle models such as ClassicJava, as demonstrated in section 4.2, and language such as PLT Scheme, as demonstrated in section 2.

In the near future, we expect to elaborate the object-oriented level in several directions. Most importantly, the object-oriented level currently assumes a functional creation mechanism for objects. While this works well for the PLT Scheme tracer, it cannot cope with Java's complex object instantiation. Conversely, the framework does not support a destroy action. To understand its addition fully, we intend to explore how an unsafe object-oriented language such as C++ can be modeled with a revised and extended framework. Furthermore, object-oriented languages often provide a modicum of reflection, which is a way of turning information about class hierarchies into objects. Last but not least we are interested in accommodating language constructs such as mixins [13] and traits [33], which require some amount of class construction at run-time.

Concerning the implementation of the framework, we intend to scale it from the kernel of PLT Scheme's class language to the full language and ProfessorJ [17], which is an implementation of Java 1.1 in PLT Scheme. Doing so will confirm that the framework is the basis for a generally useful object-oriented tracing tool. Naturally it will require some amount of human interface work to make sure the details of an execution trace are presented on demand and properly. Once finished, the implementation will help us test the usefulness of trace-based debugging with a reasonably large user base.

8. REFERENCES

- [1] M. Abadi and L. Cardelli. *A Theory of Objects*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1996.
- [2] G. Agha, I. A. Mason, S. F. Smith, and C. L. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 7(1):1–72, 1997.
- [3] S. Bernardi, S. Donatelli, and J. Merseguer. From UML sequence diagrams and statecharts to analysable petri net models. In *WOSP '02: Proceedings of the 3rd international workshop on Software and performance*, pages 35–45, New York, NY, USA, 2002. ACM Press.
- [4] K. B. Bruce. *Foundations of Object-Oriented Languages: Types and Semantics*. MIT Press, Cambridge, MA, USA, 2002.
- [5] G. Castagna, G. Ghelli, and G. Longo. A calculus for overloaded functions with subtyping. *Inf. Comput.*, 117(1):115–135, 1995.
- [6] S. M. Cho, H. H. Kim, S. D. Cha, and D. H. Bae. A semantics of sequence diagrams. *Inf. Process. Lett.*, 84(3):125–130, 2022.
- [7] W. R. Cook. *A Denotational Semantics of Inheritance*. PhD thesis, Brown University, May 1989.
- [8] W. R. Cook and J. Palsberg. A denotational semantics of inheritance and its correctness. In *OOPSLA*, pages 433–443, 1989.
- [9] S. Drossopoulou and S. Eisenbach. Java is type safe - probably. In M. Aksit and S. Matsuoka, editors, *ECOOP*, volume 1241 of *Lecture Notes in Computer Science*, pages 389–418. Springer, 1997.
- [10] S. Ducasse, T. Gırba, and R. Wuyts. Object-oriented legacy system trace-based logic testing. In *Proceedings 10th European Conference on Software Maintenance and Reengineering (CSMR 2006)*. IEEE Computer Society Press, 2006.
- [11] R. B. Findler, C. Flanagan, M. Flatt, S. Krishnamurthi, and M. Felleisen. DrScheme: A pedagogic programming environment for Scheme. In H. Glaser, P. Hartel, and H. Kuchen, editors, *Programming Languages: Implementations, Logics, and Programs*, volume 1292 of *LNCS*, pages 369–388, Southampton, UK, September 1997. Springer.
- [12] M. Flatt. PLT MzScheme: Language manual. Technical Report PLT-1, PLT, Inc., 2005. <http://www.plt-scheme.org/software/>.
- [13] M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *Proc. 25th ACM Symposium on Principles of Programming Languages*, pages 171–183, 1998.
- [14] M. Fowler and K. Scott. *UML Distilled: Applying the Standard Object Modeling Language*. Addison-Wesley Longman Ltd., Essex, UK, 1997.
- [15] P. Graunke, R. Findler, S. Krishnamurthi, and M. Felleisen. Modeling web interactions. In *European Symposium on Programming*, pages 238–252, April 2003.
- [16] P. T. Graunke. *Web Interactions*. PhD thesis, Northeastern University, April 2003.
- [17] K. E. Gray and M. Flatt. ProfessorJ: a gradual introduction to Java through language levels. In R. Crocker and G. L. S. Jr., editors, *OOPSLA Companion*, pages 170–177. ACM, 2003.
- [18] C. A. Gunter and J. C. Mitchell, editors. *Theoretical Aspects of Object-Oriented Programming: Types, Semantics, and Language Design*. MIT Press, Cambridge, MA, USA, 1994.
- [19] J. H. Hausmann, R. Heckel, and S. Sauer. Towards dynamic meta modeling of UML extensions: An extensible semantics for UML sequence diagrams. In *HCC*, pages 80–87. IEEE Computer Society, 2001.
- [20] C. Hewitt. Viewing control structures as patterns of passing messages. *Artificial Intelligence*, 8(3):323–364, 1977.
- [21] C. Hewitt, P. Bishop, and R. Steiger. A universal modular actor formalism for artificial intelligence. In *Int. J. Conference on Artificial Intelligence*, pages

- 235–245, 1973.
- [22] S. N. Kamin. Inheritance in SMALLTALK-80: a denotational definition. In *Proc. 15th ACM Symposium on Principles of Programming Languages*, pages 80–87, 1988.
- [23] J. W. Klop. Term rewriting systems: a tutorial. *Bulletin of the EATCS*, 32:143–182, June 1987.
- [24] S. Krishnamurthi, R. B. Findler, P. Graunke, and M. Felleisen. Modeling web interactions and errors. In *Interactive Computation: The New Paradigm*. Springer Verlag, 2006 – to appear.
- [25] B. Lewis. Debugging backwards in time. http://www.lamdacs.com/debugger/AADEBUG_Mar_03.pdf, September 2003.
- [26] X. Li, Z. Liu, and J. He. A formal semantics of UML sequence diagram. In *Australian Software Engineering Conference*, pages 168–177. IEEE Computer Society, 2004.
- [27] M. S. Lund and K. Stølen. Extendable and modifiable operational semantics for UML 2.0 sequence diagrams. In *Proc. 17th Nordic Workshop on Programming Theory (NWPT'2005)*, pages 86–88. DIKU, 2005.
- [28] E. Nantajeewarawat and R. Sombatsrisomboon. On the semantics of unified modeling language diagrams using Z notation. *Int. J. Intell. Syst.*, 19(1-2):79–88, 2004.
- [29] M. Okazaki, T. Aoki, and T. Katayama. Formalizing sequence diagrams and state machines using concurrent regular expression. In *Proceedings of 2nd International Workshop on Scenarios and State Machines: Models, Algorithms, and Tools SCESM'03*, pages 74–79, 2003.
- [30] U. S. Reddy. Object as closures: abstract semantics of object oriented languages. In *Proc. ACM Conference on Lisp and Functional Programming*, pages 289–297, 1988.
- [31] T. Richner. *Recovering Behavioral Design Views: a Query-Based Approach*. PhD thesis, University of Berne, May 2002.
- [32] T. Richner and S. Ducasse. Using dynamic information for the iterative recovery of collaborations and roles. In *Proceedings of ICSM '2002 (International Conference on Software Maintenance)*, Oct. 2002.
- [33] N. Schärli, S. Ducasse, O. Nierstrasz, and A. P. Black. Traits: Composable units of behaviour. In *ECOOP 2003—Object-Oriented Programming: 17th European Conference*, volume 2743, pages 248–274. Springer-Verlag, 2003.
- [34] R. J. Walker, G. C. Murphy, J. Steinbok, and M. P. Robillard. Efficient mapping of software system traces to architectural views. In S. A. MacKay and J. H. Johnson, editors, *CASCON*, page 12. IBM, 2000.
- [35] J. Whittle. On the relationship between UML sequence diagrams and state diagrams. In *Proc. of OOPSLA 2000 Workshop: Scenario-based round-trip engineering*, pages 1–6. Tampere University of Technology, Software Systems Laboratory, October 2000. Report 20.
- [36] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Inf. Comput.*, 115(1):38–94, 1994.
- [37] F. Xia and G. S. Kane. Defining the semantics of

UML class and sequence diagrams for ensuring the consistency and executability of OO software specification.
<http://cc.ee.ntu.edu.tw/~atva03/papers/16.pdf>,
 December 2003.