

First-Class State Change in Plaid

Joshua Sunshine[†] Karl Naden[†] Sven Stork^{§†} Jonathan Aldrich[†] Éric Tanter^{‡*}

[†]Carnegie Mellon University [§]University of Coimbra

[‡]PLEIAD Lab / Computer Science Department (DCC) / University of Chile

{sunshine, kbn, svens, jonathan.aldrich}@cs.cmu.edu etanter@dcc.uchile.cl

Abstract

Objects model the world, and state is fundamental to a faithful modeling. Engineers use state machines to understand and reason about state transitions, but programming languages provide little support for building software based on state abstractions. We propose Plaid, a language in which objects are modeled not just in terms of classes, but in terms of changing abstract states. Each state may have its own representation, as well as methods that may transition the object into a new state. A formal model precisely defines the semantics of core Plaid constructs such as state transition and trait-like state composition. We evaluate Plaid through a series of examples taken from the Plaid compiler and the standard libraries of Smalltalk and Java. These examples show how Plaid can more closely model state-based designs, enhancing understandability, enhancing dynamic error checking, and providing reuse benefits.

Categories and Subject Descriptors D.2.10 [Software Engineering]: Design—Representation; D.3.2 [Programming Languages]: Language Classifications—Object-oriented languages; D.3.3 [Programming Languages]: Language Constructs and Features—State

General Terms Design, Documentation, Languages, Reliability

Keywords typestate, state-chart, plaid

1. Introduction

Object-oriented programming provides a rich environment for modeling real-world and conceptual objects within the computer. Fields capture attributes of objects, methods capture their behavior, and subtyping captures specialization

relationships among objects. A key element missing from object-oriented programming languages, however, is abstract states and conceptual state change. State change is pervasive in the natural world; as a dramatic example, consider the state transition from egg, to caterpillar, to pupae, to butterfly. Modeling systems with abstract states and transitions between them is also common in many engineering disciplines.

In computer science, state machines are important modeling concepts in the UML. Abstract states are also a critical, though often implicit, part of many library APIs—and any client using that API must be aware of those states to use the API correctly. For example, a file may be in the open or closed state. In the open state, one may read or write to a file, or one may close it, which causes a *state transition* to the closed state. In the closed state, the only permitted operation is to (re-)open the file.

Files provide a simple example of abstract states, but there are many more. Streams may be open or closed, iterators may have elements available or not, collections may be empty or not¹. State spaces may be complex: In `ResultSet` from the Java JDBC library, we found 33 unique states dealing with different combinations of openness, direction, random access, insertions, etc [7]. States are also common: a recent study of protocols in Java suggests that almost three times as many types define protocols as define type parameters [3]. They also cause significant pain: for instance, in a study of problems developers experienced when using the ASP.NET framework, 3/4 of the issues identified involved temporal constraints such as the state of the framework in various callback functions [16]. All this raises a natural question: why not support abstract states in programming languages?

We previously proposed *Typestate-Oriented Programming* as a new programming paradigm in which programs are made up of dynamically created objects, each object has a typestate that is changeable, and each typestate has an interface, representation, and behavior [1]. The term typestate refers to a static abstract state checking methodology proposed by Strom and Yemini [25]; this paper focuses on a dynamically-

* Partially funded by FONDECYT Project 1110051

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

OOPSLA'11, October 22–27, 2011, Portland, Oregon, USA.
Copyright © 2011 ACM 978-1-4503-0940-0/11/10...\$10.00

¹ E.g. in Java, the cause of an exception can only be set once.

typed setting, and so we will use the terms (*abstract*) *state* and *protocol* in place of *typestate* to avoid confusion.

A programming language with abstract states can have many benefits. First, in the case of stateful abstractions, the code will more clearly reflect the intended design. This in turn will make state constraints more salient to developers who need to be aware of them. If state constraints are implicitly enforced by the object model, there is no need to code up explicit checks; thus code implementing states can be more concise. Explicit state models raise the level of error messages; instead of (perhaps) silently corrupting a data structure when an inappropriate method is called, the runtime can signal an error that that method is unavailable in the current state. Finally, explicit modeling of states also exposes new concepts for widespread reuse; candidates may include open/closed resources or positioning (beginning, middle, end) of streams.

Contribution. The contribution of this paper is the concrete design and evaluation of Plaid, an object-oriented programming language that incorporates first-class state change as well as trait-like state composition. Plaid has been implemented, and has proven effective for writing a diverse set of small and medium-sized (up to 10kLOC) programs, including a self-hosted compiler. For the purposes of this paper, Plaid is dynamically typed, though there are plans to add a gradual type system following recent work on gradual typestates [29].

The most interesting aspects of Plaid’s design come from the intersection of state change with support for a trait-like model of composition [12]. Central goals of the language design include supporting the primary state modeling constructs from statecharts [14], as well as flexible code reuse. Our design includes a hierarchical state space, so that the open state of a stream can be refined into `within` and `eof` substates indicating whether there is data left to be processed. Handling real designs in a modular way requires support for multi-dimensional state spaces, as in `and-states` from [14]; an example is a separate dimension of a stream’s state indicating whether the stream has been marked with a location or not. Modularity further requires reasoning about dimensions separately; for example, the `mark()` method should affect the marked state dimension but it should not affect whether the stream is at `eof`. Dimensions also delineate natural points of reuse; we would like to specify them separately and combine them using a trait-like composition operator.

We position Plaid relative to earlier work in the next section. Plaid’s design is described by example in section 3. This section validates our design, using a number of carefully chosen examples to concretely illustrate how Plaid provides the potential benefits described above. We also discuss our prototype implementation of Plaid, targeting the JVM.

In order to be precise about Plaid’s semantics, Section 4 provides a formal model that includes the semantics of all of Plaid’s major features. Section 5 describes how the surface Plaid language is elaborated into the core formal

model. The paper concludes with a discussion of ongoing work on Plaid, together with an argument that the concrete benefits validated by example lead to higher-level benefits in software development and evolution.

2. Background and Related Work

Plaid’s state constructs are inspired and guided by state modeling approaches such as Harel’s statecharts [14]. Other modeling approaches include Pernici’s Objects with Roles Model [20], which models objects using a set of roles, each of which can be in one of several abstract states.

Strom and Yemini proposed *typestate* as a compiler-checkable abstraction of the states of a data structure [25]. The Fugue system was the first to integrate *typestates* with an object-oriented programming language [10]. Bierhoff et al. later observed that the complexity of protocols such as the one defined by the JDBC `ResultSet` interface requires rich state modeling constructs like those proposed by Harel [7]. This paper considers a dynamically-typed setting, so we do not discuss static checkers further.

State-dependent behavior can be encoded using the State design pattern [13]. However, this pattern is less direct than the language support we propose, and it does not help with ensuring that a client only uses operations that are available in the current state.

Dynamic languages such as Self [28] provide the ability to add and remove methods, as supported by Plaid’s state change operator. Changing a delegation slot in Self can also be used to simulate state change, as can the `become` method in Smalltalk [18]. We believe that Plaid’s more structured and more declarative constructs for state modeling have advantages in terms of error checking, succinctness, and clear expression of design compared to these encodings. Plaid’s prototype-based object model is also inspired by Self’s.

Prior State-Based Languages. The Actor model [15] treats states in a first-class way, using the current state of an actor to define the response to messages in a concurrent setting.

Taivalsaari extended class-based languages with explicit definitions of logical states (modes), each with its own set of operations and corresponding implementations [27]. Plaid’s object model differs in providing explicit state transitions (rather than implicit ones determined by fields) and in allowing different fields in different states.

The Ferret language [8] provides multiple classification, in which objects can be classified in one of several states in each of multiple dimensions. Ferret attaches dimensions to classes, not other states, so dimensions cannot come and go with state changes (unlike in Plaid and Statecharts).

A number of CAD tools such as iLogic Rhapsody or IBM/Rational Rose Real-Time support a programming model based even more directly on Statecharts [14]; such models benefit from many rich state modeling features but lack the dynamism of object-oriented systems. Recently Sterkin

proposed embedding the principal features of Statecharts as a library within Groovy, providing a smoother integration with objects [23]. Our approach focuses on adding states to object-oriented languages, rather than libraries.

Other researchers have explored adding a class change primitive to statically-typed languages [4, 6, 11]. These systems, however, do not support the richness of state models (e.g. and-states) as provided in Statecharts and in Plaid.

Schaerli et al. proposed traits [12] as a composition mechanism that avoids some of the semantic ambiguities of multiple inheritance. Schaerli’s traits did not have fields, but Plaid follows prior designs [5] to add them. Like some other recent work [9, 22], Plaid does not have the flattening property, in which the composition structure of traits is compiled away and does not affect the semantics of the resulting program. We lose the simplicity of flattening but gain the ability to model structured state spaces more directly, as described below.

An initial sketch of the Plaid language design was presented earlier [1] as an instance of the Typestate-Oriented Programming paradigm. While we recap the motivation and concept of the language from this earlier work, that paper described an unimplemented language, and neither defined the language semantics nor investigated the modeling of complex state spaces, which are the key contributions of this paper. In an earlier 4-page workshop paper, we explored the need for a modular state change operator that affects only one dimension of state change at a time [2]; this paper gives the semantics for a concrete solution to that problem. Other recent work has begun to explore a gradual, permission-based type system for Plaid [29].

3. Language

In this section we will introduce Plaid by example. These examples serve the dual purpose of explaining the language and validating the concrete benefits of Plaid.

3.1 Basics of State Change

Object protocols are rules dictating the ordering of method calls on objects. The concrete state of an object with a protocol can be abstracted into a finite number of abstract states and the object transitions dynamically between these abstract states. Therefore, clients must be aware of the abstract states in order to use the object correctly.

Most programming languages provide no direct support for protocols. Instead, protocols are encoded in the language using some combination of the state design pattern [13], conditional tests on fields, and other indirect mechanisms. In Plaid, protocols are supported directly with states, which are like classes in Java, with the crucial distinction that an object’s state changes as the object evolves.

Consider the state space of files, the canonical protocol example [1], shown in Figure 1. Some files are open and some are closed. We close an open file by calling the `close` method and open a closed file by calling the `open` method.




Figure 1. State space of File.

```

1 state File {
2   val filename;
3 }
4 state OpenFile case of File = {
5   val filePtr;
6   method read() { ... }
7   method close() { this <- ClosedFile; }
8 }
9 state ClosedFile case of File {
10  method open() { this <- OpenFile; }
11 }
  
```

Listing 1. File states in Plaid

One cannot open an open file so the open file state does not include the `close` method. Similarly, one cannot read a closed file so the closed file state does not include the `read` method.

The state space of files can be encoded cleanly in Plaid as shown in Listing 1. The `state` keyword is used to define a state. The `File` state contains the fields and methods that are common between open and closed files. In this case, only the `filename` is shared. Fields are declared with the `val` keyword.

`OpenFile` and `ClosedFile` define the methods and fields that are specific to open and closed states. Both are substates of `File`. Specialization is declared with the `case of` keyword. In addition, `case of` implies orthogonality: files can either be open or closed, not both. Methods are defined with the `method` keyword. Open files have a `read` method, a file pointer field which is presumably used by the `read` method to read the file, and a `close` method. Closed files have the `open` method.

The `open` and `close` method bodies contain the most novel bit of syntax. An object referred to by a variable `x` can be changed to state `S` by writing `x <- S`. In the `open` method we transition the receiver, referred to as in Java by the keyword `this`, to the open state by writing `this <- OpenFile`.

An example file client is shown in Listing 2. The `readClosedFile` method takes a file as an argument, opens it, reads from it, closes it, and returns the value read from the file. All of the method calls are valid if a closed file is passed to the method. If an open file is passed instead the `open` method call will fail. The library writers do not need to write any special error handling code to handle this condition

```

1 method readClosedFile(f) {
2   f.open();
3   val x = f.read();
4   f.close();
5   x; //return
6 }

```

Listing 2. File client in Plaid




Figure 2. Butterfly life-cycle.

like they would in Java. This has the concrete benefit that Plaid code for the equivalent design is smaller.

In most programming languages, fields of an object are often null in certain abstract states. For example, Java files might contain a null `filePtr` when the file is closed. Null pointers are a frequent cause of runtime errors and their cause can be difficult to diagnose. For these reasons, Tony Hoare recently called null pointers a “billion dollar mistake,” and we have not repeated this mistake in Plaid.

Plaid objects are always consistent: in other languages a programmer might forget to check the state before performing an operation and perform the operation on an object in the wrong state. Similarly, the operation might fail, but with a less specific error message. For example, if a client calls the `read` method, implemented in Java without error handling, on a closed file, Java might throw a `NullPointerException` for a null dereference of `filePtr`.

3.2 State Transitions

The file state space is a complete directed graph, every pair of states is connected in both directions by an edge. Other kinds of objects have incomplete state spaces. Consider the life-cycle of a butterfly, which is illustrated by the state-space in Figure 2. A butterfly egg hatches to a caterpillar, but it cannot ‘un-hatch’. Similarly, a butterfly never transitions directly from a caterpillar to an imago, it always transforms to a chrysalis first.

To preserve the integrity of incomplete protocols, only the method receiver (`this`), can be the target of a state change operation. If Plaid did not have this restriction it would be trivial for programmers to inadvertently violate a protocol. Consider: `val x = new Egg; x<-Caterpillar; x<-Egg`. This illegal Plaid code violates the protocol by restoring a caterpillar to an egg. Instead, in legal Plaid code, methods defined in the butterfly states perform all of the state transitions.

3.3 Dimensions of State Change

Many objects in the real world are not as simple as files or butterflies. Some objects are composed of multiple states, particularly when objects are built up from reusable components. These components may change their state independently, or orthogonally. For example, cars have both gears and brakes and when the car shifts gears it has no effect on the brakes. States that change independently are in different *dimensions*. State dimensions in programming languages were introduced in [7].

More concretely, let us say a stream is in state unmarked in dimension *markable*, and state within in dimension *position*. If the object changes to state marked, also in dimension *markable*, it will lose all of the fields and methods defined in unmarked (such as `mark`), gain those in marked (such as `reset`), and keep those in *within* (such as `read`).

The full power of Plaid comes when component states are themselves composed of multiple states. In such a setting the component states are gained and lost along with their parents. Many of this kind of deep hierarchies exist in the wild [3]. For example, in the Java Database Connectivity library, the `ResultSet` interface is composed from a combination of 33 states, four levels of nesting, and eight dimensions. A slightly simplified schematic of the state space is shown in Figure 3.

The features of the language just described correspond directly to the ‘hierarchical-states’, ‘and-states’ and ‘or-states’ proposed by Harel in his seminal state-chart paper [14]. Hierarchical-states are states that are composed of other states. And-states are states that both must be present in an object—separate dimensions that are modeled using `with` composition in Plaid. Finally, or-states are states in the same dimension, and therefore only one can be present in an object—a state that is a case of another state. These features are the fundamental building blocks of the Harel state chart formalism (which forms the basis for UML state diagrams), and are naturally encoded in Plaid exactly in the manner we just described.

In the `ResultSet` diagram (Figure 3), or-states are separated by white space. For example, `Open` and `Closed` are states in one dimension, `ForwardOnly` and `Scrollable` are in another. Hierarchical-states are indicated by nesting of the state rectangles. For example, `Scrolling` is a child of `Open` and `Begin` of `Scrolling`. The names of states with children, like `Open` appear outside the state-rectangle, and the names of simple states without children, like `Inserted`, appear inside the state-rectangle. Finally, and-states are separated into *orthogonal regions* by dotted lines, so `Direction` and `Status` are and-states.

There is a natural one-to-one correspondence between the state rectangles in the diagram and the state declarations in Plaid code. A subset of the declarations for `ResultSet` states are shown in Listing 3. The or-states are all declared to be cases of their dimensions. For example, `ForwardOnly` and `Scrollable` are cases of the `Direction` dimension.

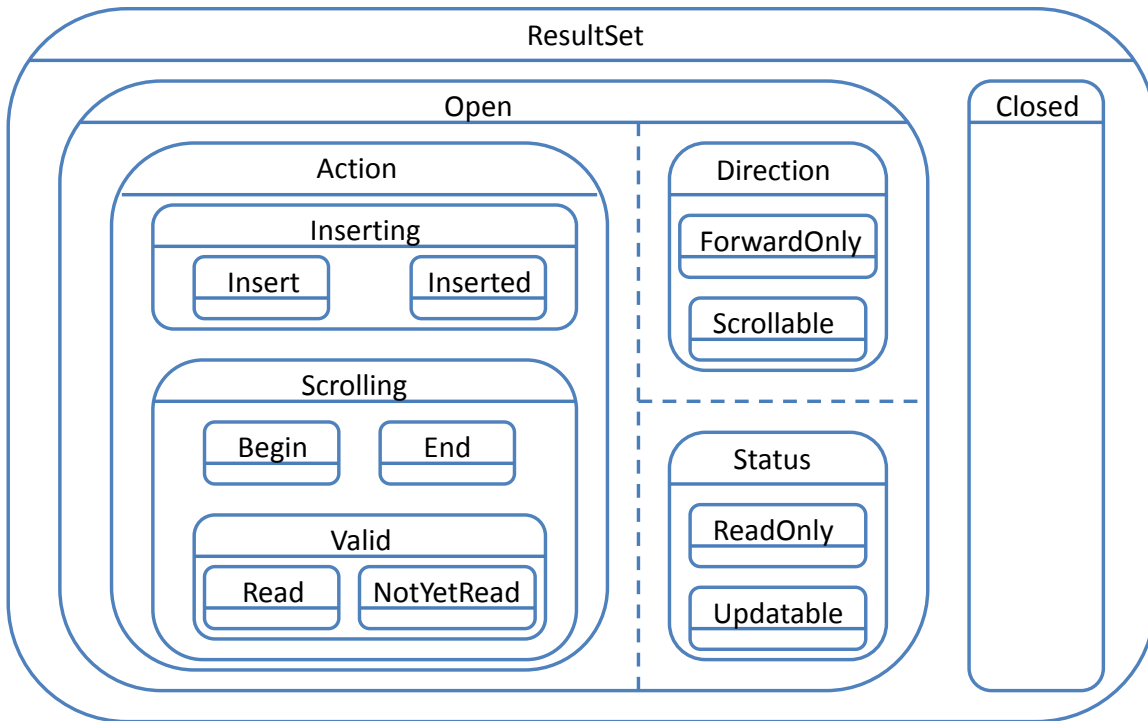


Figure 3. ResultSet state-chart.

```

1 state Open case of ResultSet =
2   Direction with Status with Action;
3 state Direction;
4 state ForwardOnly case of Direction;
5 state Scrollable case of Direction;
6 state Status;
7 state ReadOnly case of Status;
8 state Updatable case of Status;
9 state Action;
10 state Scrolling case of Action;
11 state Inserting case of Action;
12 state Insert case of Inserting;
13 state Inserted case of Inserting;
14 ...
15
16 val myResultSet = new Open @ ForwardOnly
17   with Updatable with Insert;

```

Listing 3. ResultSet state declarations and instantiation

The dimensions are themselves states in which case their or-states will inherit all of the dimension's fields and members. Sometimes, however the state is a *pure dimension* and does not contain members. In this case the state only serves to ensure that or-states do not appear together.

The and-states nested with `Open`, are declared using `with` together into the `myResultSet` state. Any object in the `Open` state is also in the `Direction`, `Status`, and `Action` states. Often `ResultSet` objects will be instantiated with children of the three dimensions `Direction`, `Status`, and `Action`.

For example, at the end of Figure 3, `myResultSet` is assigned to an open object in the `ForwardOnly`, `Updatable` and `Insert` states. This object will contain the methods and fields from `Insert`, `Inserting`, `Action`, `Updatable`, `Status`, `Forwardonly`, `Direction`, `Open` and `ResultSet`. If we were to change the state of `myResultSet` to `Inserted` by calling a method that does so for us the object would have all of the same states, with the exception that `Insert` will be replaced with `Inserted`. This is because `Insert` and `Inserted` are or-states from the same dimension. When we close the object, we lose not only the `Open` state but all of the states nested inside it. We are left only with `Closed` and `ResultSet`.

The `@` operator is syntactic sugar that allows an initializer to conveniently choose nested sub-states. The `myResultSet` initializer in Figure 3 is desugared to the following code:

```

1 var myResultSet = new Open;
2 myResultSet <- ForwardOnly with Updatable
3   with Insert;

```

First, an `Open` object is created. Then the object is changed to specializations of the three dimensions using the state change operator. Notice that the left side of the state-change operator is not `this` in the de-sugared code which violates the restriction discussed in Section 3.2. This is okay, because the restriction only applies to Plaid source which in this case uses the `@` operator.

In this example the reader can see that the Plaid code closely reflects the design embodied in the state chart. The stateful design is salient in the state declarations. Since

the mapping between the code and the state chart is so clear, a programmer reading the state declarations can easily understand the relationship between the states. In fact, our group has built a tool to automatically extract a state chart from Plural², a typestate checker for annotated Java code, and although we have not built such a tool for Plaid, the language design clearly enables it. A second potential benefit is that code for each state can be given separately in the appropriate state declaration, potentially permitting more fine-grained reuse across multiple implementations of the `ResultSet` interface.

3.4 State members

As we mentioned in the introduction, Plaid combines state change with support for a trait-like model of composition [12]. We now illustrate a particularly novel feature of Plaid, namely, state members. States can have other states as members, and these state members can be customized upon composition. This allows for consistent state update, in presence of composite states.

We illustrate state members and their benefits through a Plaid version of a `ReadWriteStream` adapted from [12], which is in turn adapted from the Smalltalk standard library. The Plaid components mirror the trait components, except in our version the methods of a single trait are sometimes divided across multiple states.

The `Position` state represents the position of the pointer into a stream or collection. It has a very limited interface which therefore makes it easy to reuse throughout an input-output and collection library. The code for `Position` is shown in Listing 4. `Position` declares two abstract methods for setting the position, a reference to the underlying collection (`vector`), constant fields for minimum and maximum position, and a variable field for the current position³.

Interestingly, `Position` contains two state members, one for the end-state and one for the not-end-state. The state members are initialized to `NotEnd` and `End`, also defined in Listing 4. These states are sub-states of `Position`, as specified by the `case of` declarations. They implement the abstract methods of `Position`. In addition, `NotEnd` has an additional method `nextPosition`, reflecting the fact that in that state, the position can be advanced. This method increments the current position, tests if the current position is at or past the maximum position, and transitions the receiver to the end state if the position is at the end. Similarly, `setToStart` in `End` transitions the receiver back to the not-end state.

The crucial part in this example is that the state transitions do not explicitly reference a specific target state, but rather reference the state members of `Position`. For instance, `nextPosition` in `NotEnd` transitions this to

```

1 state Position {
2   state notEndState = NotEnd;
3   state endState = End;
4   method setToEnd();
5   method setToStart();
6   val vector, minPos, maxPos;
7   var currPos;
8 }
9
10 state NotEnd case of Position {
11   method setToEnd() {
12     this.currPos = this.maxPos;
13     this <- this.endState;
14   }
15   method setToStart() {
16     this.currPos = this.minPos;
17   }
18   method nextPosition() {
19     this.currPos++;
20     if (this.currPos >= this.maxPos) {
21       this <- this.endState;
22     }
23   }
24 }
25
26 state End case of Position {
27   method setToEnd() { /* no op */
28   method setToStart() {
29     this.currPos = this.minPos;
30     this <- this.notEndState;
31   }
32 }

```

Listing 4. Position code.

```

1 state Reader { }
2
3 state Reading case of Reader {
4   method read() {
5     val ret = this.vector.get(this.currPos);
6     this.nextPosition();
7   }
8 }
9
10 state ReadEnd case of Reader { }
11
12 state ReadStream = Position {
13   val notEndState = Reading with NotEnd;
14   val endState = ReadEnd with End;
15 } with Reader;

```

Listing 5. ReadStream code.

this.`endState`, not `End`. This allows for consistent and flexible reuse, composition, and extension of states, as illustrated hereafter.

² <http://code.google.com/p/pluralism/>

³ Abstract methods are indicated by eliding the method body; constant fields are declared with `val`, and variable fields with `var`.

```

1 state ReadWriteStream = Position {
2   val notEndState =
3     Writing with Reading with NotEnd;
4   val endState =
5     WriteEnd with ReadEnd with End;
6 } with Reader with Writer;

```

Listing 6. ReadWriteStream code.

Consider the code for a `ReadStream`, as shown in Listing 5. The `ReadStream` definition includes a pure dimension, `Reader`. This dimension has two children `Reading` and `ReadEnd`, which correspond to the `ReadStream` in the not-end-state and the end-state, respectively. In the not-end-state, the `ReadStream` can read, and therefore `Reading` defines the `read` method. This method reads from the underlying collection at the current position and advances the position.

The `ReadStream` is composed from the two dimensions `Position` and `Reader`. `ReadStream` specializes `NotEnd` by overriding the two state members in `Position`. The state members in `ReadStream` are composed from two states, one from each dimension of `ReadStream`. Therefore, when the methods in `Position` and its children change state, they will change *both* dimensions of `ReadStream`. For example, when `nextPosition` advances the stream to the end, the `ReadStream` object composed of `Reading` with `NotEnd` will change to a `ReadEnd` with `End`.

Initializing a `ReadStream` requires two-phase initialization like for `ResultSet`. In particular, the code to create a `ReadStream` `x` that is not at the end is `val x = new ReadStream; x<-this.notEndState;` Here again, transitioning `x` to the state member `notEndState` ensures that the consistent composition of actual states is used.

Since the `Reader` dimension has the same structure as the `Position` dimension it is natural for transitions in `Position` to change `Reader` as well. In this example, there is no code in the `Reader` states that enacts the state change. Instead, the `Reader` dimension relies on the `Position` dimension to perform state changes. The state members in this example allow for this kind of dimensional reuse without extensive glue code⁴. The only code required to reuse the dimension is the specialization of state members in `ReadStream`.

We now illustrate a further step of consistent composition of states with the definition of `ReadWriteStream` in Listing 6. The definition uses a new dimension, `Writer`, with two substates `Writing` and `WriteEnd`, defined in the same manner as the `Reader` states.

This `ReadWriteStream` reuses code from all three dimensions with very little effort. The `ReadWriteStream` is the

⁴ State members also have a more traditional purposes. State members, like all states, can be used to create objects. They allow us to encode ML-style structures and functors. These abstraction mechanisms can be very powerful, especially in a typed version of Plaid. However, these purposes are not novel to Plaid so we do not focus on these here.

natural extension of `ReadStream`. The state members are composed from all three dimensions. The state transitions in a `ReadWriteStream` object will change all three dimension at once.

The `ReadWriteStream` example demonstrates both the power of a trait-like composition model and its novel extension to states. We reuse `ReadStream` and `WriteStream` with little effort, as we could achieve in a language with traits. In addition, we have a new unit of reuse, the `Position` dimension, which is shared with two other dimensions. This reuse eliminates duplicate code, and helps avoid bugs. Both the `Reader` and `Writer` of a `ReadWriteStream` are in the end-state or not-end state. Because the dimension is reused we can guarantee that no programmer will err and end up with an object in an inconsistent state like `WriteEnd` with `Reading`.

One important note is that the `Writer` and `Reader` contain no members in common, and therefore no conflict arises. Plaid requires explicit conflict resolution at the point of composition. This conflict resolution will be described in Section 4.2.

3.5 Validation

The introduction claims four concrete benefits of Plaid: code closely reflects design, programs are concise, error checking is implicit, and new opportunities for reuse. These benefits were illustrated in the examples in this section and they were discussed while describing the examples. We summarize the case here for emphasis. We then reflect on our experience writing mid-sized programs in Plaid, in diverse domains.

3.5.1 Concrete benefits

Code reflects design. Designs with stateful abstractions are clearly reflected in Plaid code. This is clear in of all three examples in this section. The implementation of the file, result set, and read-write streams all match their designs. Arbitrarily complex state-charts can be encoded in Plaid with the simple rules described alongside the result set example. Each abstract state maps to its own state in code, so the design of the abstraction and its protocol as a whole is highly *salient* in the code.

Concise programs. Since state constraints are implicitly enforced by the object model, none of our examples included any error checking code. The implementation are therefore smaller.

Error Prevention. Plaid's explicit state models make error checking more consistent, because the programmer cannot forget to check state constraints when a method is called. The level of abstraction of error messages is also thereby raised: when an inappropriate method is called, instead of triggering an internal run-time exception such as a null pointer, or (what is worse) silently corrupting data, the runtime can signal an

Project	Lines of Code	# Files
CodeGenerator	1205	24
AeminiumCodeGen	2610	8
Typechecker	4196	55
ASTtranslator	9506	107
PlaidApps	528	21
Standard Library	372	18
TestCompiler	2811	96
TestTypechecker	363	9
Total	21591	338

Table 1. Plaid code written for eight projects.

error that a particular method is unavailable in the current state. Also, we have shown how state members can be used to enforce consistency of multiple dimensions of state at once.

Reuse. Plaid provides new reuse opportunities. Some state machines are used in many objects. For instance, the `Position` dimension was reused in both read and write streams, and it could also be reused in many IO and Collection libraries. Open and closed resources like the `File` and `ResultSet` are also very common.

3.5.2 Applicability to diverse domains

In order to gain practical experience with the language and experiment with typestate-oriented programming beyond small examples, we have written several mid-sized programs in Plaid. These programs further demonstrate the expressiveness of Plaid in a diverse set of domains including compilation, input-output, GUIs, and web. They are all available for download from the Plaid repository⁵. In total, we have written 22KLOC across 338 files. A breakdown of our implementations is in Table 1. We call out items of particular interest here.

Compiler. Plaid is self-hosting; the `CodeGenerator` project compiles Plaid code into Java source. Plaid code can easily use Java libraries and many of our examples are implemented that way. In a sister project [24], we have implemented a separate compiler for parallel-by-default code, which is the `AeminiumCodeGen` project. We are currently working on a Plaid typechecker; the implementation is the `Typechecker` project. All these projects are supported by AST transformations performed by the `ASTtranslator` project.

GUI Library. GUI libraries often impose state constraints on their clients. We implemented Plaid wrappers for a few key Java Swing classes, including `Window`, `Pane`, and `Canvas` abstractions. We use states to enforce proper initialization of these abstractions. In particular, windows should have some contents added, otherwise they are created with size zero. Furthermore, windows are `Hidden` until `show()` is called, then they become `Visible`. Panes should also have contents

⁵ <http://code.google.com/p/plaid-lang/>

Obj Val	ov	$::=$	$mv \mid dv \mid mv \times ov \mid dv \times ov$
Dim Val	dv	$::=$	$tag\{ov\} \mid tag\{ov\} <: dv$
Mbr Val	mv	$::=$	$method\ m(\bar{x})\{e\} \mid$ $val\ n = v$
ObjExp	oe	$::=$	$me \times oe \mid de \times oe \mid e \times oe \mid$ $me \mid de \mid e$
Dim Exp	de	$::=$	$dv \mid tag\{oe\} \mid tag\{oe\} <: de \mid$ $e \mid e\{\bar{to}\}$
Mbr Exp	me	$::=$	$mv \mid val\ f \triangleright x = e \mid$ $recstate\{\ val\ s \triangleright x = proto\ sd\}$
State Decl	sd	$::=$	$freshtag\{oe\} <: de \mid$ $freshtag\{oe\} \mid oe$
Trait Op	to	$::=$	$\setminus n \mid n \rightarrow n' \mid me \mid (tagOf\ e).me$
Val	v	$::=$	$\ell \mid ov \mid proto\ oe \mid fn(\bar{x}) \Rightarrow e$
Exp	e	$::=$	$x \mid v \mid let\ x = e\ in\ e \mid$ $e(\bar{e}) \mid e.m(\bar{e}) \mid e.n \mid$ $e \leftarrow e \mid e \leftarrow e \mid new\ e \mid$ $match(e)\{\bar{c}\} \mid$ $freeze\ e \mid recstate\{\overline{mv}\}\#1$
Case	c	$::=$	$case(tagOf\ e)\{e\} \mid default\{e\}$

Figure 4. Internal Syntax

added. Both panes and canvases must be assigned a parent window, and canvases should be given a preferred size. Our library is not comprehensive, but it is sufficient to build demonstration applications—in our case, a Turing machine that uses Plaid’s states to represent the finite state control, the marks on the tape, and the illusion of an infinite tape. Both the windowing library and Turing demonstration application are in the `PlaidApps` project.

Miscellaneous The `Plaidapps` project includes the examples discussed earlier and a small web server and workflow engine. The Plaid standard library includes integers, rationals, strings, options, and standard control (e.g. `if`) and looping (e.g. `for`, `while`) structures. Finally, two testing projects include a number of smaller tests and examples.

4. Semantics

In this section we present the formal definition of the Plaid language and give it a precise semantics. At its core, Plaid is an object system with first-class generators and functions. Individual generators can be combined and specialized using composition and operators inspired by traits [12], instantiated to create objects, or used to specify the abstract state the object should change to. We start by describing the syntax and object model of a core language, which is intended to be simpler than Plaid source code yet be capable of representing all of the major semantic elements of Plaid. Then we discuss the execution semantics of the core language.

4.1 Core Syntax

The syntax of the internal representation of Plaid is given in Figure 4. In these definitions, x ranges over bound variables, while members of objects are represented by f , m , and s , which respectively range over fields, methods, and state members. We use n to represent any kind of object members when we do not distinguish between them. Abstract states are represented using *tags* which are generated as needed. We will introduce each syntactic category in turn, describing its purpose and motivations.

4.1.1 Expression Syntax

Plaid contains the standard expressions found in object systems, including object creation through `new`, field selection, and method calls. Because Plaid also has first-class functions, we include standard function definition and application as well. For sequential expressions, we include `let` bindings and bound variable references.

The rest of the expression forms are related to Plaid's encoding of abstract states and the transitions between them:

Changing state. The Plaid core has two state change operators. `←` represents a state update and only removes portions of the receiving object that are mutually exclusive with the incoming states. For completeness and flexibility, Plaid also includes a state replacement operator, `←←`, which wipes the receiving object clean before adding the incoming states, much like an in-place `new` operation. One could imagine using this operator in a situation where an object needed to be in a particular state and no other states. This cannot be guaranteed by the state update operator because state update leaves dimensions unrelated to the updating state alone.

Unlike the source language, Plaid's core does not require the target of a state change operator to be *this*. This makes the core simpler and more flexible since the restriction can be enforced at the source level.

proto values. First class instance generators are provided by proto expressions. These are values which can be stored in fields and passed as parameters. During a well-formed execution, the target of new expressions and the right-hand side of state change expressions will evaluate to a proto value. This is because they encapsulate object expressions, *oe*, which are uninitialized objects. The state change and new expressions cause the initialization steps specified by the object expression wrapped in the proto to be evaluated for use in creating a new object or changing the state of an existing one.

State expressions. To allow states to be chosen dynamically at runtime, we include several expression forms that can evaluate to a proto. As they are values, standard deference or bound variables could result in proto expressions. Because most states included in protocols must be defined with (mutual) recursion, proto values represented source-declared

states are wrapped into a `recstate`. A particular proto can be selected from the `recstate` as from a standard record.

The `freeze` expression is a more novel way to get a proto. It takes the object and wraps it up in a proto allowing new instances to be generated from it. As an example of the use of `freeze`, consider the `myResultSet` value defined in Listing 3. Say we wanted to do some extra initialization of the `ResultSet` before using it and that over the course of a program we would create the same `ResultSet` over and over. To avoid needing to do the same initialization repeatedly, one could freeze the object the fully initialized object and then instantiated it each time a new `ResultSet` of this form was needed. `freeze` has already been used in the Plaid compiler to more cleanly support certain initialization paradigms, such as the transformation to `let-normal` form, where strings of `let` bindings must be concatenated together.

Matching tags. Finally, the `match` construct allows pattern matching based on tags. Each case tests the target object against the `tagOf` another expression. This expression is expected to evaluate to a proto value with a single outer tag which is grabbed by `tagOf` and compared with the tags of the target object. If the object contains the tag, the corresponding case is executed. Cases are evaluated in order.

An example of the use of `match` comes from the Plaid standard library. Plaid's syntax does not include control structures. Instead, `if` and `while` are encoded as functions that make use of `match`. The states `True` and `False` are each defined as a case of `Boolean`. Thus, the `if` function determines whether or not to evaluate the body based on whether the object returned by the condition matches the `True` tag.

4.1.2 Object Value Syntax

Plaid objects are collections of tags representing the states that the object is in along with fields and methods that provide the representation and operations of those states. In order to implement the desired semantics these object must be organized to formally encode the relationships between tags and members that the semantics depend on. In particular, we need to represent the following relationships between the abstract states that the tags represent:

1. *Superstates:* An object in state, S , which is defined to be a case of a superstate, T , must also be in state T . For instance, an object in the `NotEnd` state defined in figure 4 is also in the `Position` state.
2. *Or-states:* Distinct cases of a given state, such as the `OpenFile` and `ClosedFile` case of `File`, cannot exist together in an object.
3. *And-states:* Both objects and states can be defined as a composite of other states. For example, the `Open` state from Listing 3 is defined in terms of states `Direction`, `Status`, and `Action`. Objects in the composite state are considered to be in each of the component states as well.

4. *Defining states*: Members must be associated with the state that declares them so that they can be removed from an object when their defining state is removed.

To formalize these relationships, *objects values* are organized as hierarchical collections of *dimensions*, which contain tags for the state and all of its transitive super states, and *members*.

Object values. The basic component of an object is an object value, *ov*, which is a list of dimension values, *dv*, and member values, *mv*. They are used to represent both top-level objects and the dimensions and members that define a given state (see dimension values below). The γ operator that separates each element of the list represents composition. Object values encode and-states by allowing two dimensions to coexist together inside the definition of a state. For instance, the object value that defines a `ReadStream` would have two composed dimensions, one for the `Position` dimension, and the other for the `Reader` dimension.

Dimension values. Dimension values, $tag\{ov\}[\prec: dv]$, encode the structure of a state and its super states. They are represented by a tag, *tag*, which is a unique name for the most specialized state from the dimension. Associated with the tag is an object value which represents the collection of members that the state defines along with any other dimensions that make up the and-states of the state. A dimension value may optionally contain another dimension value encoding the superstate relationship.

By containing the representation of a given states super-states, dimension values give us a way to encode the or-state relationship as well. Two states that are the case of the same superstate would be encoded as separate dimensions with the same state at the root of the dimension. Because the tags in the dimensions partially overlap, by restricting tags to appear only once in a given object value, we can ensure that no or-states can coexist in a single object.

Concretely, we would represent an instantiated `Open` state from Listing 3 as

```
Open{Direction  $\gamma$  Status  $\gamma$  Action}  $\prec$ : ResultSet.
```

Here the most specific state of the represented dimension value is `Open`. This state is defined based on the three states `Direction`, `Status`, and `Action` (defining object values not shown), and specializes the `ResultSet` state, which it was defined as a case of.

A dimension is also Plaid’s version of a trait. Multiple inheritance is achieved by allowing multiple dimensions to be composed in an object value as well as in the object values associated with the tags of a dimension. The hierarchical nature of Plaid’s dimension prevent us from using all of the trait mechanisms for solving the problems of multiple inheritance. In particular, a multiple inheritance system must deal with the case when one class inherits from two classes that share a (transitive) parent. This situation is challenging because it is non-obvious how to inherit members from the

common grandparent. This problem is commonly referred to as the *diamond problem* [19], because of the shape of the inheritance hierarchy diagram. The original traits proposal [12] *flattens*⁶ composed traits and forces any conflicts between method names to be explicitly resolved (field were not allowed in traits). However, as Plaid’s semantics depend on members being related to the tag they are defined in, we cannot use flattening. Instead, Plaid prevents the diamond problem by preventing or-states from coexisting, thereby preventing the same tag and member definition from appearing more than once (following Malayeri’s no-diamonds rule [19]). Plaid’s solution follows recent extensions of traits including [5, 9, 22]. Like Plaid, these system support traits with fields and work in a variety of object models including those that, like Plaid, add hierarchy and do not enforce the flattening property. As with the original trait proposal, all name conflicts across dimensions must be explicitly resolved in Plaid via the trait operators described below.

Member values. A member value is either a method, with a set of arguments and a body, or a field, `val f`, bound to a value, *v*. The member is said to be defined in the state represented by its immediately enclosing tag. As a concrete example, an object in the `ClosedFile` state described in Listing 1 would be represented formally as

```
ClosedFile{method close(){e}}  
   $\prec$ : File{val filename = v}
```

This indicates that the object is in both the `ClosedFile` and the `File` states, one of which is a substate of the other, and each of which defines a single member.

4.1.3 Uninitialized Object Syntax

Plaid has corresponding syntax for uninitialized objects organized into *object expressions*, *dimension expressions*, and *member expressions*. When compared to their value counterparts, they share the same structure but contain expressions which are not yet values. In this section, we discuss the places where execution can occur in these forms and the motivation behind them.

Object expressions. Object expressions, *oe*, are made up of the composition of dimension expressions, member expressions, as well as raw expressions. The purpose of unevaluated expressions in dimension and member expressions will be explained below. Raw expressions as components of object expressions allow part of an uninitialized object to be determined at the time of initialization. These expressions evaluate to proto values which are then incorporated into the initializing object. This provides for Plaid’s implementation of dynamic trait composition by allowing portions of the object to be selected at runtime.

⁶The flattening property from [12] states that members of an are treated equally regardless of what trait they were defined in.

Dimension Expressions Dimension expression can contain unexecuted expressions in the object expression associated with the most specific tag as well as in tags up the hierarchy if they exist. Dimension expressions may also have associated trait operations, to , which need to be evaluated. Trait operations allow standard manipulations such as renaming, $n \rightarrow n'$, and removal, $\setminus n$. Note that these operate on the *whole* dimension, renaming or removing all members of the specified name defined directly in tags in the hierarchy (not including nested dimensions). This allows the changes to be preserved by state change in the dimension as we will see below.

Members can also be added or replaced⁷. By default, they are (re)placed in the most specific tag of the dimension expression. However, in cases where members need to be added or replaced in a particular tag, they can be qualified by a particular tag, specified as with tags in case statements by `tagOf` another expression. The redefinition of `Position.EndState` for the `ReadStream` in Listing 5 is an example of using qualified trait operations. This mechanism is important in Plaid because of the hierarchical nature of Plaid’s object model and when and how member definitions are removed during state change.

Member expressions. Only fields can be member expressions, me , as methods do not have any initialization code. On the other hand, fields can be defined with initialization expressions that require evaluation as a part of object creation or update. In order to allow fields to refer to the initialized value of previous fields in the same state, field expressions define an internal bound variable in addition to their external name (this is a standard approach from [21], chapter 8). Fields are also generated by state declarations. Since the definitions of related states, such as the `OpenFile` and `ClosedFile` from Listing 1, are typically recursive, the initialization of state members occurs in a `recstate` binding.

State members are also special in that when an uninitialized object containing state members is initialized, new tags may need to be generated. The `proto` expression encapsulates uninitialized objects as discussed above. Normally they contain object expressions, but when appearing in a `recstate`, they contain state declarations, sd which may contain the `freshTag` operations that generates a new tag when executed, resulting in an object expression. This feature means that new tags are generated for states defined inside states each time the outer state is instantiated. Because these tags can then be used to pattern match on objects, this allows Plaid to implement ML-style generative functors⁸. Functors have

⁷The semantics defined here do not allow fields and states in trait operations to refer to other trait operation members. The formalism could be extended to support this, mirroring the case for declarations in states

⁸Generative functors, in contrast to applicative functors, generate new abstract types for each application of the functor. This impacts pattern matching when using these generated types in a similar way as pattern matching on freshly generated tags in Plaid.

Heap	H	$::=$	$[\ell \rightsquigarrow ov], H \mid \cdot$
Eval	E	$::=$	$[\] \mid \text{let } x = E \text{ in } e \mid E(\bar{e}) \mid v(\bar{v}, E, \bar{e}) \mid E.m(\bar{e}) \mid v.m(\bar{v}, E, \bar{e}) \mid E.f \mid E \leftarrow e \mid v \leftarrow E \mid v \leftarrow \text{proto } E \mid E \leftarrow e \mid v \leftarrow E \mid v \leftarrow \text{proto } E \mid \text{new } E \mid \text{new proto } E \mid \text{match}(E)\{\bar{c}\} \mid \text{match}(v)\{\text{case}(\text{tagOf } E) \{e\}, \bar{c}\} \mid \text{freeze } E \mid ov \ \gamma \ E \mid O \ \gamma \ oe \mid O$
Obj	O	$::=$	$\text{val } n \triangleright x = E \mid \text{tag}\{oe\} <: E \mid \text{tag}\{E\} \mid \text{tag}\{E\} <: dv \mid E\{\bar{to}\} \mid dv\{\bar{to}, \text{val } n = E, \bar{to}\} \mid dv\{\bar{to}, (\text{tagOf } e).(\text{val } n = E), \bar{to}\} \mid dv\{\bar{to}, (\text{tagOf } E).mv, \bar{to}\}$

Figure 5. Contexts

well recognized modularity benefits that we do not discuss here.

4.2 Dynamic Semantics

We now introduce the dynamic semantics of Plaid. We formalize the execution using a small step operational semantics. The basic evaluation judgment has the form $e@H \mapsto e'@H'$ and is read “expression e with heap H evaluates to expression e' in heap H' ”. We define a similar judgment $oe@H \mapsto oe'@H'$ for the evaluation of object expressions. In this section, we will define the form of the heap and the invariants that we maintain on it. We will also discuss the Plaid-specific evaluation rules, in particular those that use ancillary judgments for implementing state change. As state change is at the core of Plaid’s design and is the most complicated we go into depth about the motivation and design of the rules that implement it. Finally, we describe object initialization and trait operations that may be involved.

4.2.1 Heap

A heap, H , is a mapping from locations, ℓ , to object values. We place additional well-formedness requirements on all object values stored in the heap. These restrictions prevent ambiguities from multiple inheritance.

Tag uniqueness. We require that all well-formed object values have no duplicate tags. As alluded to above, this property ensures that an object is not in two cases of a single or-state at the same time. This is because the tags representing two mutually exclusive or-states must come from the same dimension and thus must have at the least the root tag of the dimension in common. It also prevents the diamond problem of multiple inheritance by ensuring that a particular member definition does not appear multiple times in a single object. This invariant is encoded in the helper judgment `uniqueTags` also defined in Figure 10.

Member uniqueness. Even though a given definition for a member cannot appear more than once, it is still possible

$e@H \mapsto e@H$	
$\frac{}{\text{let } x = v \text{ in } e@H \mapsto e[v/x]@H} \text{E-LET}$	$\frac{ \{\bar{x}\} = \{\bar{v}\} }{(\text{fn } (\bar{x}) \Rightarrow e)(\bar{v})@H \mapsto e[v/x]@H} \text{E-APP}$
$\frac{H[\ell] = ov \quad \text{lookup}(m, ov) = (\text{method } m(\bar{x})\{e\}) \quad \{\bar{x}\} = \{\bar{v}\} }{\ell.m(\bar{v})@H \mapsto e[\ell/\text{this}][v/x]@H} \text{E-CALL}$	$\frac{H[\ell] = ov \quad \text{lookup}(f, ov) = (\text{val } f = v)}{\ell.f@H \mapsto v@H} \text{E-FIELD}$
$\frac{H[\ell] = ov_1 \quad \text{uniqueTags}(ov_2) \quad ov_1 \leftarrow ov_2 \Rightarrow ov_3 \quad \text{uniqueMembers}(ov_3)}{\ell \leftarrow \text{proto}(ov_2)@H \mapsto \text{void}@H[\ell \rightsquigarrow ov_3]} \text{E-SU}$	$\frac{\text{uniqueTags}(ov) \quad \text{uniqueMembers}(ov)}{\ell \leftarrow \text{proto}(ov)@H \mapsto \text{void}@H[\ell \rightsquigarrow ov]} \text{E-REPLACE}$
$\frac{\ell \notin H \quad \text{uniqueTags}(ov) \quad \text{uniqueMembers}(ov)}{\text{new } (\text{proto } ov)@H \mapsto \ell@H[\ell \rightsquigarrow ov]} \text{E-NEW}$	
$\frac{de = \text{tag}\{oe\}\{<: de'\} \quad \text{tag} \notin \text{tags}(H[\ell])}{\text{match}(\ell)\{\text{case } (\text{tagOf } \text{proto } de)\{e\}, \bar{C}\}@H \mapsto \text{match}(\ell)\{\bar{C}\}@H} \text{E-CASENOMATCH}$	
$\frac{de = \text{tag}\{oe\}\{<: de'\} \quad \text{tag} \in \text{tags}(H[\ell])}{\text{match}(\ell)\{\text{case } (\text{tagOf } \text{proto } de)\{e\}, \bar{C}\}@H \mapsto e@H} \text{E-CASEMATCH}$	
$\frac{}{\text{match}(\ell)\{\text{default}\{e\}, \bar{C}\}@H \mapsto e@H} \text{E-CASEDEFAULT}$	$\frac{H[\ell] = ov}{\text{freeze } \ell@H \mapsto \text{proto } ov@H} \text{E-FREEZE}$
$\frac{1 = 1_s \quad oe = oe_{1_s} [\text{recstate}\{\text{val } s_i = \text{proto } oe_i\} \# 1_i / s_i]}{\text{recstate}\{\text{val } s_i = \text{proto } oe_i\} \# 1@H \mapsto \text{proto } oe@H} \text{E-RECSTATESELECT}$	

Figure 6. Expression Evaluation

that multiple tags define members with the same name. To prevent ambiguities in this case we require that all members of an object are provided by exactly one dimension. Because the hierarchy of dimensions gives us a natural way to choose the visible definition (the one from the most specific tag in the dimension) we allow a single name to be defined directly in multiple tags from a single dimension. Formally, two tags are in the same dimension if one is a transitive case of the other. This relaxation of classical traits allows, for instance, a common super state to define a default behavior for a method which can be overridden by (some of) its substates. The judgment `uniqueMembers` defined in Figure 10 captures this requirement. It uses the judgments $mv :: \text{tag}.x@ov$, which states that member value mv from tag tag defines name x in object value ov , and $\text{tag} << \text{tag}'@ov$ which asserts the property that tag tag is a transitive subtag of tag' in object value ov . Based on these helper judgments, an object value has unique members if whenever we find the same member defined in two tags, then one of these tags is a transitive subtag

of the other. We prove that evaluation preserves member and tag uniqueness in the companion technical report [26].

Member lookup. As an object can contain multiple members with the same name, we need an unambiguous way to choose which one is visible. The `lookup` function also in Figure 10 defines this logic. When multiple definitions are found, we know by `uniqueMembers` that they all come from the same dimension. Since the tags of a dimension form a total order, we know that one of tags defining the member will be a transitive subtag of all other tags defining the member. The definition from this most specific tag is the one returned by `lookup`.

4.2.2 Expressions

The evaluation rules for expressions in Plaid are given in Figure 6. We only list only computation rules here, defining congruence rules using evaluation contexts shown in Figure 5. In these, each expression with a subexpression that requires evaluation defines a hole, $[]$, into which any expression can

be placed. Evaluation proceeds by using the computation rule that evaluates the expression in the hole.

Standard rules. The computation rules for the evaluation of the expressions from general object systems and the lambda calculus are almost all completely standard in our system. These include the rules E-LET, E-APP, E-CALL, and E-FIELD for let expressions, application, method calls and field dereferences respectively. One note is that member selection during calls and dereferences use the `lookup` judgment described above. We also use standard record evaluation rules when selecting a label from a `recstate` (E-RECSTATESELECT).

Match. Plaid uses a first-match semantics, so that we find the first case clause whose tag matches the target object. We find the tag to match against by grabbing the most specific tag (`tagOf`) from a dimension expression wrapped in a proto value. Note that in this case the dimension expression is not evaluated since we are only interested in the tag. If the tag is found in the target object, the code for this case is evaluated (E-CASEMATCH); otherwise, execution proceeds to the next case (E-CASENOTMATCH). Default cases are always executed and terminate the match if reached (E-CASEDEFAULT). Evaluation gets stuck if no matching case is found.

Freezing. To freeze a location in the heap (E-FREEZE), we simply pull the object value from the heap and wrap it in a proto expression.

Manipulating objects in the heap. The state change operators and new cause objects in the heap to be changed or allocated. Because we only allow object values to appear in the heap, we must first initialize the object that will be used to alter the heap by reducing it to an object value. Evaluation is mostly handled by the evaluation contexts: first the expression representing the object is reduced to a proto value and then the object expression wrapped in the proto is evaluated down to an object value. An important design decision in Plaid was to run the initializers for all members of an object expression. This happens despite the fact that not all members may end up in the object (see the explanation of state update below). In particular, any effectful initializers will always be run and update the wider context. We experimented with other possible semantics but decided that a clear and unambiguous rule for when initializers were run (always) was better than a flexible but complicated one. Furthermore, we consider it good Plaid style to avoid the use of effectful initializers and instead use other design techniques, such as factory methods, when effectful operations are required as a part of object initialization.

Once the initialization code in the proto has been run, the resulting object value can be used to update the heap. In the case of new and state replacement (\leftarrow) expressions it is clear what the object value that is inserted into the heap will be. `new` allocates a new location on the heap and maps it

```
1 val rs = Open {
2   Inserted <: Inserting <: Action,
3   Scrollable <: Direction,
4   Updatable <: Status }
5 <: ResultSet
```

Listing 7. Open, Inserted, Scrollable, Updatable ResultSet

to the resulting object value. State replacement replaces the mapping of the target location on the heap with the updating object value. Since we know the precise form of the object value that is being inserted into the heap, in order to maintain the heap invariants on object values, we can simply check that `uniqueTags` and `uniqueMembers` both hold on the new object value as done in the rules E-REPLACE and E-NEW. On the other hand, the semantics of updating an object on the heap using state update are much more complicated, and so we devote the next section to a discussion of its design and proof that they maintain the necessary invariants.

4.2.3 State Update

At the core of the rule E-SU which updates the heap with the result of a state update is the state update judgment, $ov \leftarrow ov \Rightarrow ov$, which is described in Figure 7. The judgment takes two object values and determines the resulting object value when the *target* object on the left side of the arrow is changed to the state given by the *update* object from the right side. The semantics of this judgement are the most complicated and important part of Plaid’s dynamic semantics. Thus, before describing the semantics given by the rules, we step back and give a high-level overview of the desired behavior. We then define some general properties and assumptions of the judgment before describing the rules themselves.

Design considerations. Our goal is that the design of the state change judgment should match the semantics of stateful abstractions as modeled by state charts and similar tools. Thus, a state update should transition a target object from its current set of abstract states to a possibly new set of abstract states as specified by the update object. To do this, we need to formalize this intuition in terms of object values.

Update dimensions. Our first task is to determine which abstract state the update object is changing. That is, which dimensions of the target object need to be updated? Consider the object value (without members) of an `Open ResultSet` in the `Inserted`, `Scrollable`, `Updatable` state, stored in `val rs` as depicted in Listing 7. What should happen if we update `rs` to the `ReadOnly` state?

$$rs \leftarrow \text{ReadOnly} <: \text{Status}$$

While there are clearly matches between tags in the target and update objects, since the tags are nested inside the `Open` tag of the target object, it is not clear that they should be updated.

$$\boxed{ov \leftarrow ov \Rightarrow ov}$$

$$\frac{ov_t \leftarrow ov \Rightarrow ov' \quad ov' \leftarrow ov_u \Rightarrow ov_o}{ov_t \leftarrow ov \times ov_u \Rightarrow ov_o} \text{SU-LIST}$$

$$\frac{}{ov_t \leftarrow mv_u \Rightarrow ov_t \times mv_u} \text{SU-MV}$$

$$\frac{\text{tags}(ov_t) \cap \text{tags}(dv_u) = \emptyset \quad \text{uniqueTags}(dv_u)}{ov_t \leftarrow dv_u \Rightarrow ov_t \times dv_u} \text{SU-ADDH}$$

$$\frac{\text{tags}(dv) \cap \text{outerTags}(dv_u) \neq \emptyset \quad dv \leftarrow dv_u \Rightarrow dv_r \quad \text{tags}(ov) \cap \text{tags}(dv_r) = \emptyset}{ov \times dv \leftarrow dv_u \Rightarrow ov \times dv_r} \text{SU-MATCHDIM}$$

$$\frac{\text{outerTags}(dv_u) \cap \text{tags}(ov) \neq \emptyset \quad ov \leftarrow dv_u \Rightarrow ov_r \quad [\text{tags}(dv_u) \cap \text{tags}(dv) = \emptyset] \quad tag \notin \text{tags}(dv_u)}{tag\{ov\}[\lt;: dv] \leftarrow dv_u \Rightarrow tag\{ov_r\}[\lt;: dv]} \text{SU-MATCHINNER}$$

$$\frac{\text{outerTags}(dv_u) \cap \text{innerTags}(dv) \neq \emptyset \quad dv \leftarrow dv_u \Rightarrow dv_r \quad \text{tags}(tag\{ov\}) \cap \text{tags}(dv_u) = \emptyset}{tag\{ov\} \lt;: dv \leftarrow dv_u \Rightarrow tag\{ov\} \lt;: dv_r} \text{SU-MATCHSUPERINNER}$$

$$\frac{tag \notin \text{outerTags}(dv_u) \quad \text{outerTags}(dv_u) \cap \text{outerTags}(dv) \neq \emptyset \quad dv \leftarrow dv_u \Rightarrow dv_r}{tag\{ov\} \lt;: dv \leftarrow dv_u \Rightarrow tag\{ov\} \lt;: dv_r} \text{SU-MATCHSUPER}$$

$$\frac{dv_u = [dv_{sub}] \lt;: tag\{ov'\} \lt;: [dv_{sup}] \quad [\text{tags}(dv_{sub}) \cap \text{tags}(tag\{ov\}[\lt;: dv])] = \emptyset \quad \text{uniqueTags}(dv_{sub})}{tag\{ov\}[\lt;: dv] \leftarrow dv_u \Rightarrow [dv_{sub}] \lt;: tag\{ov\}[\lt;: dv]} \text{SU-MATCH}$$

Figure 7. State Update

However, if we think of the state update as an transition to a new abstract state, then we can see that the nesting in the target object should not matter. This state update specifies that the Status dimension should transition to the ReadOnly substate, and thus out of the Updatable state.

The converse question is does nesting matter in the other direction? In other words, can a nested state trigger a change in an abstract state? Concretely, would this state update

$$rs \leftarrow \text{Foo}\{\text{ReadOnly} \lt;: \text{Status}\}$$

result in an object in the ReadOnly state? Based on the semantics of state charts, the answer would be “no”. Our definitions of object dimension indicates that the Status dimension of the Foo state is part of the definition of Foo. Thus, it is brought along with the transition to the Foo state. The Status state is also a defining and-state of the Open state. Thus the resulting object cannot be consistent because two separate dimensions are claiming the Status state meaning there would need to be duplicate tags.

Therefore, we define the dimensions along which a state update occurs to be only those found at the top level of the object value that describes the update object. All other dimensions that are a part of the update object are considered definitions of these dimensions and do not induce transitions but are only added to the object with their enclosing state.

Dimension updates. Once we know which dimensions will be updated, we need to know what in those dimensions is changed. We first note that we can treat the transition in each dimension independently as dimensions are orthogonal by definition. Second, recall the file example from Listing 1. In this example, we stated that the filename member was shared between the OpenFile and ClosedFile states. Thus, when we transition from an ClosedFile to an OpenFile the members of the File state should remain constant. This is the semantics behind the restricted update semantics of state change described in [2]. We use and extend these semantics in a natural way to account for our hierarchical object model.

Properties of object and state update. With the intuition we have for the design, we can define some terminology that is used in the judgment itself.

Inner and outer tags. In the informal description of state change, we differentiated between dimensions and tags defined at the top level of the update object and those that appear within a top-level dimension. Figure 10 defines two judgments, outerTags and innerTags, which capture this distinction. The outerTags of an object value, ov , are all the tags which appear as the most specific tag and any of its super tags from dimensions appearing directly in ov . For example, using rs , the ResultSet object from Listing 7, $\text{outerTags}(rs) = \{\text{Open}, \text{ResultSet}\}$. Conversely, the innerTags of an object value are all of the tags defined in

dimensions that are recursively included in the definition of each of the outer tags. For example,

$$\text{innerTags}(rs) = \{\text{Inserted}, \text{Inserting}, \text{Action}, \text{Scrollable}, \text{Direction}, \text{Updatable}, \text{Status}\}$$

Unique dimension property. Given a dimension within which to transition the target object, we need to find the location of the matching dimension within the target object value. To do this, we look for the part of the object that has tags which overlap the outer tags of the update dimension. We ignore all super-tags of the matching tag in the update dimension under the assumption that these supertags will match the tags in the target. This assumption is based on the the Unique Dimension Property which states that a single unique tag can only ever appear in a single dimension. That is a tag is either has no super tags or always appears with the same supertag. While this property is not guaranteed by the syntax and semantics of the internal language, it is enforced by the elaboration from Plaid’s source syntax so we assume it in our rules.

Maintaining the uniqueTags property. Rule E-SU in Figure 6 does not check whether the object returned from the state update judgement has unique tags. Therefore the state update judgement must maintain this property. Formally: If $\text{uniqueTags}(ov_1) \wedge ov_1 \leftarrow ov_2 \Rightarrow ov_3$, then $\text{uniqueTags}(ov_3)$. A proof of this property is in the companion technical report [26].

Inference rules. With this understanding, we can describe the rules that produce the object value after a state update operation. The rules start by breaking apart the update object ov into the individual member values and dimension values and processing the state changes for each dimension or value individually (SU-LIST). This is allowed since each dimension can be treated independently. We can assume that uniqueTags holds for each dimension individually since it holds for the object as a whole. For member values (SU-MV) and dimension values for which there is no overlap between the tags of the target object and update dimension (SU-ADDH), we just compose the update object with the target object. The rest of the rules assume that there is a match between the outer tags of the update object and the tags of the target object. If that is not the case, then the evaluation gets stuck.

SU-MATCHDIM covers the case where we have found a particular dimension of the target object that contains the tags that are changing. By the unique dimension property explained above, we know that the $\text{outerTags}(dv_u)$ will not appear in ov , so it suffices to calculate the state update on just the matched dimension. To ensure that we maintain the unique tags property, we can assume that both the result of the state update and the unmatched portion of the object have

unique tags, and so it suffices to check that the tags of these two portions of the object do not intersect.

SU-MATCHINNER handles the case where there is overlap between the innerTags of the current tag and the outerTags of the update dimension. We recursively find the state update on just this matching portion and then check that the tags from the resulting object value do not intersect with the tags of the super tag, if it exists, to maintain the uniqueTags invariant.

In SU-MATCHSUPERINNER, we find that the matching dimension is defined somewhere inside of a super tag. Thus, we run state update on just the supertags. We then verify that the tags of the result are distinct from the tags of the subtag and its innerTags to maintain the uniqueTags invariant.

SU-MATCHSUPER represents the case where we have found the right dimension, but have not reached the level of the dimension where the tags overlap. The current tag of this dimension is not in the outer tags of the update dimension, but there is overlap somewhere in its super tags and so we find the updated state from that portion of the dimension. In this case, we know that the current tag will be removed with any of its nested tags, which means that we do not need to check if these tags would conflict with tags that enter the object with the update dimension to preserve uniqueTags .

The base case SU-MATCH handles the actual alteration of the target dimension. The current tag matches a specific tag in the outer tags of the update dimension, which indicates that the state update only affects states below this point in the dimension. In particular the tags below this one in the dimension in the target object are discarded, as already occurred through the SU-MATCHSUPER rule. In their place are put all the subtags of the matched tag from the incoming dimension. To make sure that we do not have duplicate tags anywhere, we only need to check that the tags added from the update dimension do not intersect with the tags that are in its new supertags.

Example. To give a specific example, consider evaluating the following state update on the object defined in Listing 7:

$$rs \leftarrow \text{ReadOnly} <: \text{Status}$$

The state updates proceeds first by finding that there is tag overlap between the incoming and target objects and a match for the `Status` tag of the incoming state nested inside the `Open` state with the SU-MATCHINNER. Next it finds the correct dimension `Updateable <: Status` using the SU-MATCHDIM rule. It discards the `Updateable` tag and recurses up the dimension in the SU-MATCHSUPER rule and finally adds the `ReadOnly` tag in its place with the SU-MATCH rule.

Reduction rule The E-SU reduction rule uses the state update judgement to determine what object value to update the target object to. The state update judgement incrementally checked that uniqueTags was maintained. It does not

$$\boxed{oe@H; T \mapsto oe@H; T}$$

$$\frac{[oe' = oe[v/f']]}{\text{val } f \triangleright f' = v[\gamma oe]@H \mapsto \text{val } f = v[\gamma oe']@H} \text{E-RECFIELD}$$

$$\frac{sd = \text{freshtag}\{oe\}[\langle: de\rangle] \quad \text{tag is fresh}}{r = \text{recstate}\{\text{val } s_d \triangleright x_d = v, \text{val } s \triangleright x = \text{proto } (tag\{oe\}[\langle: de\rangle]), \text{val } s_r \triangleright x_r = \text{proto } s_{d_r}\}[\gamma oe]}{\text{recstate}\{\text{val } s_d \triangleright x_d = v, \text{val } s \triangleright x = \text{proto } sd, \text{val } s_r \triangleright x_r = \text{proto } s_{d_r}\}[\gamma oe]@H \mapsto r@H'} \text{E-RECSTATE1}$$

$$\frac{oe'_i = oe_i[\text{recstate}\{\text{val } s_i = \text{proto } oe_i\} \# s_i/x_i] \quad [oe' = oe[\text{recstate}\{\text{val } s_i = \text{proto } oe\} \# s_i/x_i]]}{\text{recstate}\{\text{val } s_i \triangleright x_i = \text{proto } oe_i\}[\gamma oe]@H \mapsto \text{val } s_i = \text{proto } oe'_i[\gamma oe']@H} \text{E-RECSTATE2}$$

$$\frac{}{(tag\{oe\} \langle: \text{proto } de[\{\overline{to}\}]\}[\gamma oe']@H \mapsto (tag\{oe\} \langle: de[\{\overline{to}\}]\}[\gamma oe']@H) \text{E-DE}$$

$$\frac{}{\text{proto}[ov \gamma](\text{proto } oe[\{\overline{to}\}])[\gamma oe']@H \mapsto \text{proto}[ov \gamma](oe[\{\overline{to}\}])[\gamma oe']@H} \text{E-OE}$$

$$\frac{dv\{\overline{to}\} \Rightarrow dv'}{(dv\{\overline{to}\})[\gamma oe]@H \mapsto dv'[\gamma oe]@H} \text{E-TRAITOPS}$$

Figure 8. Object Evaluation

guarantee that `uniqueMembers` is satisfied and so the rule checks that the resulting object value has unique member declarations.

4.2.4 Object Evaluation

The final class of reductions that we must model is that of state expressions, including the initialization of object expressions within a proto. These rules are defined in Figure 8. Congruence rules are again taken care of by evaluation contexts from Figure 5.

- **E-RECFIELD:** When field members have been evaluated down to values, we propagate them forward into the rest of the declarations that need to be initialized by substituting the value in for the bound variable on the right of the \triangleright . This allows subsequent fields to use the values of previously declared fields during their initialization. After this propagation, we do not need to keep track of the bound variable any longer and so do not record it in the member value. Note that these semantics force us to be strict about the order in which portions of the object are initialized. In particular member declarations are initialized from left to right as specified by the evaluation contexts.
- **E-RECSTATE1:** If there are `freshtag` directives in the state declarations of a `recstate`, new tags are generated by picking a fresh `tag` not previously mentioned.
- **E-RECSTATE2:** After assigning new tags to all of the state declarations inside a `recstate`, we need to remove the `recstate` construct and convert it into a list of `val` declarations. This is done in a manner similar to the `fix` construct in the lambda calculus. Since our `recstate` is

modeled as a record, we replace all references to the inner bound variable of each of the nested state `vals` with selections of the external name from the `recstate`. We do this both inside the object expressions of each `proto` as well as in subsequent declarations. Note again that after propagation we can remove the bound variable from the `val` declaration.

- **E-DE and E-OE:** These rules state that it is possible to unwrap a proto that is nested inside another proto. This can occur when a proto is part of an object expression inside another proto (E-OE), or when a proto is in a dimension expression, which only appear in proto expressions (E-DE). In either case, if trait operations are associated with this proto, then they are retained. Execution will continue by evaluating the wrapped object expression if needed.
- **E-TRAITOPS:** This rule applies only once the all of the trait operations have been fully reduced and proceeds using the trait operations judgment defined below to produce a new dimension value.

4.2.5 Trait Operations

As with state change, we define a separate judgement for trait operations that applies once all trait operations have been fully initialized, meaning that they can all be applied atomically without reduction. The rules for initialization of trait operations are all congruence rules handled by evaluation contexts (see Figure 5). Thus, the judgement, $ov\{\overline{sp}\} \Rightarrow ov$, does not require a heap. In general, trait operations follows previous work on traits. However, Plaid's object model, unlike traditional traits models, is hierarchical. Hence, trait

$$\boxed{ov\{\overline{to}\} \Rightarrow ov}$$

$$\frac{dv\{to\} \Rightarrow dv' \quad dv'\{\overline{to}\} \Rightarrow dv''}{dv\{to, \overline{to}\} \Rightarrow dv''} \text{T-GENERAL} \qquad \frac{ov = ov'[\gamma mv'] \quad [\text{name}(mv') = x = \text{name}(mv)]}{(tag\{ov\}[\lt;: dv])\{mv\} \Rightarrow tag\{ov' \gamma mv'\}[\lt;: dv]} \text{T-MEMBER}$$

$$\frac{ov\{\setminus n\} \Rightarrow ov' \quad [dv\{\setminus n\} \Rightarrow dv']}{(tag\{ov\}[\lt;: dv])\{\setminus n\} \Rightarrow tag\{ov'\}[\lt;: dv']} \text{T-REMOVEDV} \qquad \frac{dv\{\setminus n\} \Rightarrow dv' \quad ov\{\setminus n\} \Rightarrow ov'}{(dv \gamma ov)\{\setminus n\} \Rightarrow dv' \gamma ov'} \text{T-REMOVEOV1}$$

$$\frac{\text{name}(mv) = n \quad ov\{\setminus n\} \Rightarrow ov'}{(mv \gamma ov)\{\setminus n\} \Rightarrow ov'} \text{T-REMOVEOV2} \qquad \frac{\text{name}(mv) \neq n \quad ov\{\setminus n\} \Rightarrow ov'}{(mv \gamma ov)\{\setminus n\} \Rightarrow mv \gamma ov'} \text{T-REMOVEOV3}$$

$$\frac{ov\{n \rightarrow n'\} \Rightarrow ov' \quad [dv\{n \rightarrow n'\} \Rightarrow dv']}{(tag\{ov\}[\lt;: dv])\{\setminus n\} \Rightarrow tag\{ov'\}[\lt;: dv']} \text{T-RENAMEDV} \qquad \frac{dv\{n \rightarrow n'\} \Rightarrow dv' \quad ov\{n \rightarrow n'\} \Rightarrow ov'}{(dv \gamma ov)\{n \rightarrow n'\} \Rightarrow dv' \gamma ov'} \text{T-RENAMEOV1}$$

$$\frac{\text{name}(mv) = n \quad \text{rename}(n', mv) = mv' \quad ov\{n \rightarrow n'\} \Rightarrow ov'}{(mv \gamma ov)\{n \rightarrow n'\} \Rightarrow mv' \gamma ov'} \text{T-RENAMEOV2}$$

$$\frac{\text{name}(mv) \neq n \quad ov\{n \rightarrow n'\} \Rightarrow ov'}{(mv \gamma ov)\{n \rightarrow n'\} \Rightarrow mv \gamma ov'} \text{T-RENAMEOV3} \qquad \frac{de = tag\{oe\}[\lt;: de'] \quad ov\{tag.mv\} \Rightarrow ov'}{(ov)\{\text{tagOf proto } de\}.mv\} \Rightarrow ov'} \text{T-STATEMEMBER}$$

$$\frac{tag \notin \text{tags}(dv) \quad ov\{tag.mv\} \Rightarrow ov'}{(dv \gamma ov)\{tag.mv\} \Rightarrow dv \gamma ov'} \text{T-STATEMEMBEROV1} \qquad \frac{tag \in \text{tags}(dv) \quad dv\{tag.mv\} \Rightarrow dv'}{(dv[\gamma ov])\{tag.mv\} \Rightarrow dv'[\gamma ov]} \text{T-STATEMEMBEROV2}$$

$$\frac{tag \neq tag' \quad ov\{tag'.mv\} \Rightarrow ov' \quad [dv\{tag.mv\} \Rightarrow dv']}{(tag\{ov\}[\lt;: dv])\{tag'.mv\} \Rightarrow tag\{ov'\}[\lt;: dv']} \text{T-STATEMEMBERDV1}$$

$$\frac{(tag\{ov\}[\lt;: dv])\{mv\} \Rightarrow dv'}{(tag\{ov\}[\lt;: dv])\{tag.mv\} \Rightarrow dv'} \text{T-STATEMEMBERDV2}$$

Figure 9. Trait Operations

operations other than the local member addition must take this hierarchy into account.

Local member updates are agnostic to whether the added member is already a member of the tag and simply add the new member, replacing the existing member if one exists (T-MEMBER). Updates of members in specific tags act the same, but first must recurse through the object value looking for the specified tag before performing the member update. The computation will get stuck if the tag is not found. Because each of these trait operations, as well as member renaming described below, may potentially add new members, there is the danger that the object value might no longer satisfy the `uniqueMembers` invariant. However, since the specialization must be occurring as part of object instantiation, it will be checked at the point that the object is created, so we do not make the check here.

Member removal and renaming operate on the whole object, removing or renaming instances of members with the given name throughout. This is in contrast to `lookup`, which stops at the first declaration of the member. These semantics are required in order to allow trait composition,

which includes the ability to remove members from a trait and instead provide them in another trait. This would result in a conflict if some members were left in the old dimension.

5. Elaboration

The core language defined in the previous section shares much in common with the full Plaid programming language, but there are still differences. The semantics of the full Plaid language are defined as an elaboration into the core language. Space does not permit including the full elaboration in this paper, but we present the highlights and refer the interested reader to the accompanying technical report [26].

For most expressions, the elaboration proceeds structurally, without changing the construct itself. For field bindings, we add the internal variable referred to above, and replace references to the field in later field initializers with the fresh variable. Sequences of state declarations are transformed into `recstate` blocks. Each state declaration is transformed into a `val` declaration which binds to a `proto` representing the uninitialized state, with a `freshTag` expression for generating the state's tag when the declaration is executed.

$\text{uniqueTags}(ov)$	$\text{uniqueMembers}(ov)$	$\text{lookup}(x, ov) = mv$	$dv \in ov$	$mv :: \text{tag}.x@ov$
$\text{tag} <<: \text{tag}@ov$	$\text{validTagMembers}(ov)$	$\text{rename}(n, mv) = mv$	$\text{name}(mv) = n$	
$\text{tags}(ov)$	$\text{outerTags}(ov)$	$\text{innerTags}(ov)$		

$$\frac{\text{tag} \notin \text{tags}(ov) \quad [\cup \text{tags}(dv)] \quad [\text{tags}(ov) \cap \text{tags}(dv) = \emptyset]}{\text{uniqueTags}(ov) \quad [\text{uniqueTags}(dv)]} \text{UNIQUETAGSDV}$$

$$\text{uniqueTags}(\text{tag}\{ov\}[\langle: dv])$$

$$\frac{\text{tags}(dv) \cap \text{tags}(ov) = \emptyset \quad \text{uniqueTags}(dv) \quad \text{uniqueTags}(ov)}{\text{uniqueTags}(dv \ \gamma \ ov)} \text{UNIQUETAGSOV1}$$

$$\frac{[\text{uniqueTags}(ov)]}{\text{uniqueTags}(mv[\ \gamma \ ov])} \text{UNIQUETAGSOV2} \quad \frac{mv_1 :: \text{tag}_1.x@ov \ \dots \ mv_n :: \text{tag}_n.x@ov \quad \text{tag}_i <<: \text{tag}_i@ov \ \dots \ \text{tag}_i <<: \text{tag}_i@ov}{\text{lookup}(x, ov) = mv_i} \text{LOOKUP}$$

$$\frac{\text{validTagMembers}(ov) \quad \exists n (\exists \text{tag} \ mv :: \text{tag}.n@ov \wedge \exists \text{tag}' \ mv' :: \text{tag}'.n@ov) \implies (\text{tag} <<: \text{tag}'@ov \vee \text{tag}' <<: \text{tag}@ov)}{\text{uniqueMembers}(ov)} \text{UNIQUEMEMBERS}$$

$$\frac{}{dv \in dv} \text{LEAF1} \quad \frac{dv \in ov'}{dv \in mv \ \gamma \ ov'} \text{LEAF2} \quad \frac{\text{tag} \neq \text{tag}' \quad \text{tag}\{ov\}[\langle: dv] \in dv'}{\text{tag}\{ov\}[\langle: dv] \in \text{tag}'\{ov'\}[\langle: dv'] \ \gamma \ ov'']} \text{LEAF3}$$

$$\frac{\text{tag} \neq \text{tag}' \quad \text{tag}\{ov\}[\langle: dv] \in ov'}{\text{tag}\{ov\}[\langle: dv] \in \text{tag}'\{ov'\}[\langle: dv'] \ \gamma \ ov'']} \text{LEAF4} \quad \frac{\text{tag}\{[ov_1 \ \gamma \ mv[\ \gamma \ ov_1]]\}[\langle: dv] \in ov}{mv :: \text{tag}@ov} \text{MBRINTAG}$$

$$\frac{\text{tag}\{ov'\}[\langle: dv] \in ov \quad \text{tag}' \in \text{outerTags}(\text{tag}\{ov'\}[\langle: dv])}{\text{tag} <<: \text{tag}'@ov} \text{CASEOF} \quad \frac{\text{name}(mv) \notin \text{names} \quad [\text{validTagMembers}(\text{names} \cup \text{name}(mv), ov')]}{\text{validTagMembers}(\text{names}, mv[\ \gamma \ ov'])} \text{VTM1}$$

$$\frac{\text{validTagMembers}(\emptyset, ov) \quad [\text{validTagMembers}(\emptyset, dv)] \quad [\text{validTagMembers}(\text{names}, ov')]}{\text{validTagMembers}(\text{names}, (\text{tag}\{ov\}[\langle: dv])[\ \gamma \ ov'])} \text{VTM2}$$

$$\frac{}{n = \text{name}(\text{val } n = v)} \text{NAME1} \quad \frac{}{m = \text{name}(\text{method } m(\bar{x})\{e\})} \text{NAME2} \quad \frac{}{\text{rename}(a, \text{val } n = v) = \text{val } a = v} \text{RENAME1}$$

$$\frac{}{\text{rename}(n, \text{method } m(\bar{x})\{e\}) = \text{method } n(\bar{x})\{e\}} \text{RENAME2} \quad \frac{}{\text{tags}(ov) = \text{innerTags}(ov) \cup \text{outerTags}(ov)} \text{TAGS}$$

$$\frac{}{\text{outerTags}(\text{tag}\{ov\}[\langle: dv]) = \{\text{tag}\} \cup \text{outerTags}(dv)} \text{OUTERDV}$$

$$\frac{}{\text{outerTags}(dv[\ \gamma \ ov]) = \text{outerTags}(dv) \cup \text{outerTags}(ov)} \text{OUTEROV1}$$

$$\frac{}{\text{outerTags}(mv[\ \gamma \ ov]) = \emptyset \cup \text{outerTags}(ov)} \text{OUTEROV2}$$

$$\frac{}{\text{innerTags}(\text{tag}\{ov\}[\langle: dv]) = \text{tags}(ov) \cup \text{innerTags}(dv)} \text{INNERDV}$$

$$\frac{}{\text{innerTags}(dv[\ \gamma \ ov]) = \text{innerTags}(dv) \cup \text{innerTags}(ov)} \text{INNEROV1}$$

$$\frac{}{\text{innerTags}(mv[\ \gamma \ ov]) = \emptyset \cup \text{innerTags}(ov)} \text{INNEROV2}$$

Figure 10. Helper Judgements

Declarations	D	$::=$	$SD \mid \text{method } m(\bar{x})\{SE\} \mid \text{val } f = SE$
State Decl.	SD	$::=$	$\text{val } s = S \mid \text{state } s = S \mid \text{state } s \text{ case of } s\{\overline{TO}\} = S$
States	S	$::=$	$\text{freeze}(SE) \mid \{\overline{D}\} \mid s\{\overline{T}\} \mid S \text{ with } S \mid SE.s \mid s$
Trait Ops	TO	$::=$	$\backslash n \mid n \rightarrow n' \mid \text{val } f = SE \mid \text{val } s = S \mid \text{val } s.f = SE \mid \text{val } s.t = S \mid \text{method } m(\bar{x})\{SE\} \mid \text{method } s.m(\bar{x})\{SE\}$
Expression	SE	$::=$	$x \mid \text{let } x = SE \text{ in } SE \mid SE.f \mid SE(\overline{SE}) \mid SE.m(\overline{SE}) \mid SE \leftarrow S \mid SE \leftarrow S \mid \text{new } S \mid \text{match}(SE)\{\overline{C}\} \mid$
Case	C	$::=$	$\text{case } SE.s \{SE\} \mid \text{case } s \{SE\} \mid \text{default } \{SE\}$
Compil. Unit	CU	$::=$	\overline{D}

Figure 11. Source Syntax

Our formal semantics defines all of the Plaid language except for module linking and cross language binding. Module linking currently follows the Java standard, including packages, imports, and a classpath for loading elements. Plaid primitives are defined using Java classes and methods, which can be directly accessed in Plaid via their fully-qualified Java names. Details of both of these aspects of Plaid are discussed in more detail in the Plaid language definition [17].

6. Discussion and Future Work

The primary contribution of the Plaid language is providing a way for programmers to express state machine abstractions directly in the source code of their programs. Plaid supports the major state modeling features of Statecharts, including state hierarchy, or-states, and and-states. The explicit representation of states makes the design more salient in the code, enhancing programmer understanding. For example, the separation of members into different abstract states helps programmers quickly learn what operations are available in each state. In the future, visualization tools that leverage explicit state constructs to automatically generate statecharts from Plaid code could provide even greater benefits.

Plaid has the potential to make code more reliable. Not only do explicit states help programmers understand libraries better, avoiding errors in the first place; the runtime will also verify that the libraries are used correctly according to their state abstractions. Even a “method not available in this state” error is better than a silent corruption, but in future work, we believe we can leverage explicit states to do much better. For example, a state-related error message could be paired with a suggestion about what methods could be called in order to move the object into a correct state.

Plaid’s trait-like state composition model provides a way of reusing not just fields and methods, but state abstractions. This additional layer of reuse has the potential to reduce redundancy in code and specifications, while enhancing developer productivity. The confidence that comes with the error checking in Plaid’s state model may also help developers to evolve and refactor software with greater confidence.

In future work, we plan to build more programs with Plaid in order to investigate the possible benefits outlined above. We are also developing a gradual type system that can complement Plaid’s dynamic state checking with static checking, where desired by programmers [29]. We believe Plaid demonstrates a new kind of language, and we are excited to explore the consequences that language may entail.

Acknowledgments

We thank Nels Beckman and Robert Bocchino for their work on the semantics of Plaid; Manuel Mohr, Mark Hahnenberg, Aparup Banerjee, Matthew Rodriguez, and Fuyao Zhao for their work on the Plaid compiler; and the PLAID group for their helpful feedback and suggestions. This research was supported by DARPA grant #HR00110710019. Sven Stork is supported by the Portuguese Research Agency FCT, through a scholarship (SFRH/BD/33522/2008). Joshua Sunshine was supported by the Department of Defense (DoD) through the National Defense Science and Engineering Graduate Fellowship (NDSEG) Program.

References

- [1] J. Aldrich, J. Sunshine, D. Saini, and Z. Sparks. Tpestate-oriented programming. In *Proc. Onward*, 2009.
- [2] J. Aldrich, K. Naden, and É. Tanter. Modular composition and state update in Plaid. In *Proc. Mechanisms for Specialization, Generalization and Inheritance*, MASPEGHI, 2010.
- [3] N. E. Beckman, D. Kim, and J. Aldrich. An empirical study of object protocols in the wild. In *European Conference on Object-Oriented Programming 2011*.
- [4] A. Bejleri, J. Aldrich, and K. Bierhoff. Ego: Controlling the Power of Simplicity. In *Proc. Foundations of Object-Oriented Languages*, 2006.
- [5] A. Bergel, S. Ducasse, O. Nierstrasz, and R. Wuyts. Stateful traits and their formalization. *Journal of Computer Languages, Systems and Structures*, 34(2):83–108, 2008.
- [6] L. Bettini, S. Capecchi, and F. Damiani. A Mechanism for Flexible Dynamic Trait Replacement. In *Proc. Formal Techniques for Java-like Programs*, 2009.
- [7] K. Bierhoff and J. Aldrich. Lightweight object specification with tpestates. In *Proc. Foundations of Software Engineering*, 2005.
- [8] B. Bloom, P. Keyser, I. Simmonds, and M. Wegman. Ferret: Programming language support for multiple dynamic classification. *Computer Languages, Systems and Structures*, 35(3): 306 – 321, 2009.

- [9] T. V. Cutsem, A. Bergel, S. Ducasse, and W. D. Meuter. Adding state and visibility control to traits using lexical nesting. In *Proc. European Conference on Object-Oriented Programming*, 2009.
- [10] R. Deline and M. Fahndrich. Typestates for Objects. In *Proc. European Conference on Object-Oriented Programming*, 2004.
- [11] S. Drossopoulou, F. Damiani, M. Dezani-Ciancaglini, and P. Giannini. Fickle: Dynamic Object Re-classification. In *Proc. European Conference on Object-Oriented Programming*, 2001.
- [12] S. Ducasse, O. Nierstrasz, N. Schärli, R. Wuyts, and A. Black. Traits: A mechanism for fine-grained reuse. *ACM Transactions on Programming Languages and Systems*, 28(2):331–388, 2006.
- [13] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design patterns: elements of reusable object-oriented software*. Addison-Wesley, 1995.
- [14] D. Harel. Statecharts: A visual formalism for complex systems. *Science of computer programming*, 8(3):231–274, 1987.
- [15] C. Hewitt, P. Bishop, and R. Steiger. A Universal Modular Actor Formalism for Artificial Intelligence. In *Proc. International Joint Conference on Artificial Intelligence*, 1973.
- [16] C. Jaspan. Proper plugin protocols: Cost-effective verification of frameworks. Technical Report CMU-ISR-11-101, Institute for Software Research, Carnegie Mellon University, April 2011. Thesis Proposal, originally accepted April 2010.
- [17] Jonathan Aldrich. *The Plaid Language: Dynamic Core Specification*, 2010. <http://plaid-lang.googlecode.com/hg/docs/spec/current/current.pdf>.
- [18] A. C. Kay. The Early History of Smalltalk. *SIGPLAN Notices*, 28(3), 1993.
- [19] D. Malayeri and J. Aldrich. CZ: multiple inheritance without diamonds. *Proc. Object-Oriented Programming, Systems, Languages, and Applications*, 2009.
- [20] B. Pernici. Objects with Roles. In *Proc. Conference on Office Information Systems*, 1990.
- [21] B. C. Pierce. *Advanced Topics in Types and Programming Languages*. MIT Press, 2005.
- [22] J. Reppy and A. Turon. Metaprogramming with traits. In *Proc. European Conference on Object-Oriented Programming*, 2007.
- [23] A. Sterkin. State[chart]-Oriented Programming. In *Proc. Multiparadigm Programming with Object-Oriented Languages*, 2008.
- [24] S. Stork, P. Marques, and J. Aldrich. Concurrency by default: using permissions to express dataflow in stateful programs. In *Proceeding of the 24th ACM SIGPLAN conference companion on Object oriented programming systems languages and applications*, pages 933–940. ACM, 2009.
- [25] R. Strom and S. Yemini. Typestate: A programming language concept for enhancing software reliability. *IEEE Transactions on Software Engineering*, 12(1):157–171, 1986.
- [26] J. Sunshine, K. Naden, S. Stork, J. Aldrich, and É. Tanter. First-class state change in plaid. Technical Report CMU-ISR-11-114, Institute for Software Research, Carnegie Mellon University, October 2011.
- [27] A. Taivalsaari. Object-Oriented Programming with Modes. *Journal of Object-Oriented Programming*, 6(3):25–32, 1993.
- [28] D. Ungar and R. B. Smith. Self: The Power of Simplicity. In *Proc. Object-Oriented Programming, Systems, Languages, and Applications*, 1987.
- [29] R. Wolff, R. Garcia, É. Tanter, and J. Aldrich. Gradual typestate. In *European Conference on Object-Oriented Programming 2011*.