

Developing a Formal Semantics for Babelsberg: A Step-by-Step Approach

Tim Felgentreff, Todd Millstein, and Alan Borning

VPRI Technical Report TR-2014-002
September 2014

1 Introduction

Babelsberg [4] is a family of object constraint languages, with current instances being Babelsberg/R (a Ruby extension), Babelsberg/JS [5] (a Javascript extension), and Babelsberg/S (a Squeak extension). The Babelsberg design has evolved alongside its implementations, driven in part by practical considerations about the expectations of programmers familiar with the underlying host language. This fact, along with the complexities of integrating objects, state, and constraints, have led to a number of the semantic choices being muddled with implementation specifics. There have also been a number of long-standing, confusing issues with respect to constraints and object identity, how to represent assignment, and the appropriate restrictions on expressions that define constraints. In an effort to understand these better and to provide a complete design that instances of Babelsberg can implement, we give here a formal operational semantics for Babelsberg.

We've found it helpful to approach the problem incrementally, first devising a formal semantics for a very simple constraint imperative language, and then building up to a full object constraint language. In this memo we present the semantics in that fashion as well. The languages are as follows:

- Babelsberg/PrimitiveTypes (the only datatypes are booleans plus a set of primitive types, such as integers, reals, and strings). The concrete variant we use in informal examples is Babelsberg/Reals.
- Babelsberg/Records (Babelsberg with immutable records, along with primitive types).
- Babelsberg/UID (Babelsberg with mutable records that live on the heap and so have an identity, can be aliased, etc., as well as primitive types).
- Babelsberg/Objects (Babelsberg with mutable objects, classes, methods, messages, inheritance, and object-oriented constraint definitions).

In each case, we first provide an informal discussion and examples, and then the formal semantics. The primary audience for this admittedly long tech report is first ourselves — we have clarified many aspects of the language as a result of working out the formalism — and second, other researchers who are interested in the details. Subsequently we plan to write a standard-length paper for a programming language conference that should be of more general interest, based on this material. (This subsequent paper will focus just on Babelsberg/Objects, perhaps with a brief introductory example using Babelsberg/Reals, and will also include proofs of some of the interesting properties of the language.)

2 Motivation

Our formal semantics is intended to provide a complete semantics of Babelsberg that can be used to inform practical implementations of the language. It is meant to be as simple as possible, while still encompassing the major design decisions needed to guide language implementers. Because Babelsberg is a design to provide object-constraint programming in an object-oriented host language, the semantics omits some constructs such as exception handling for constraint solver failures and syntactic sugar that are intended to be inherited from the host language. The semantics instead focuses on the expression of standard object-oriented constructs that need to be modified to support the Babelsberg design.

An overarching design goal is that in the absence of constraints, Babelsberg should be a standard object-oriented language. We have thus resisted the temptation to add other interesting features; instead we aim to make the smallest possible changes while still accommodating constraints in a clean and powerful way.

We require that our design support a useful and expressive language for constraints, integrated with the host object oriented language in a clean way. Beyond this, however, a heuristic for the design is that when there is a choice, we favor simplicity over power (at least power to do interesting operations in the language that have nonetheless not yet proven themselves useful in practice, in particular with respect to object identity and type). The constraint imperative language Kaleidoscope [6, 7, 8, 9], particularly the early versions, was arguably too complex in part because it was too powerful in interesting ways of just this kind, making it difficult to understand what the result of a program might be and also difficult to implement efficiently. Trying to follow this heuristic also makes the design more independent of the particular host language and solvers used in implementing it, because only a small set of basic operations have to be adapted.

Some significant clarifications and simplifications of Babelsberg that have come out of this work on formalizing the language are:

- A clearer understanding of the interactions among constraints on identity, types, and values.
- The addition of structural compatibility checks to tame the power of the solver with respect to changing object structure and type to satisfy constraints.
- The addition of *value classes*. (Instances of value classes are immutable objects for which object identity is not significant.) Value classes play a key role in giving a clear specification of the requirements for expressions that define constraints (in particular that they must be side effect free), while still supporting useful programming idioms.
- A set of restrictions on identity constraints that make it easier to reason about object identity and type. In particular, any change to the identity of the object referred to by a variable flows clearly from a single assignment statement, and is deterministic — there are never multiple correct solutions to the identity constraints. This also implies that any change to the type of a variable must similarly flow from a single assignment statement and be deterministic.

3 Constraints

The semantics for each of these languages includes a step in which we assert that some set of values for variables is a correct solution to a collection of hard and soft constraints. There is a constraint solver that is a black box as far as the rest of the formal semantics is concerned, and that handles all the issues regarding finding a solution, dealing with conflicting soft constraints, and so forth. The solver should be sound but may be incomplete (i.e., it should never return an incorrect answer, but might respond that the set of constraints is too difficult for it to determine whether or not there is a solution).

We use the semantics for hard and soft constraints presented in [3]. An earlier paper on Babelsberg [4] has a description of the relevant theory as well. The knowledge about how to handle hard and soft constraints is left entirely up to the solver and doesn't enter into the formal semantics, so the discussion in the rest of this section is intended to help the reader understand the informal examples — it isn't relevant to the formal semantics.

The solver should find a single best solution — if there are multiple solutions, the solver is free to pick any one of them. (Providing answers rather than solutions, i.e., results such as $10 \leq x \leq 20$ rather than a single value for x , and backtracking among multiple answers, as available in for example constraint logic programming [11], is left for future work — see Appendix A.6.)

The way we trade off conflicting soft constraints is defined by a *comparator*. The solver encodes the comparator being used, making it irrelevant to the formal semantics — however, in presenting the informal program examples we'll sometimes need to specify which is used, and for this reason we include a brief discussion of comparators.

The two most relevant comparators are locally-predicate-better (LPB) and weighted-sum-better (WSB). Locally-predicate-better only cares whether a constraint is satisfied or not, not how far off the value is from the desired one. Any Pareto-optimal solution is acceptable. For example, a solution that satisfies one weak constraint A and violates three weak constraints B, C, and D is OK, as long as there isn't a solution that satisfies both A and some additional constraint, even if there is another solution that satisfies B, C, and D but not A. The DeltaBlue solver [10] finds LPB solutions. Weighted-sum-better considers the weighted sum of the errors of constraints with a given soft priority, and picks a solution that minimizes the sum. If there is more than one solution and there are additional lower priority constraints, we then consider the lower-priority ones to winnow down the possible solutions, priority by priority. For this comparator we need an error in satisfying a constraint, which should be 0 iff the constraint is satisfied. Cassowary [2] finds WSB solutions.

Here are two examples. (These are described from the point of view of the declarative theory of hard and soft constraints, not with respect to how an actual solver can find that solution.)

```
required  x + y = 10
strong    x = 8
weak      y = 0
```

The required constraint has an infinite number of solutions. When we winnow these down to solutions that satisfy the strong constraint, there is only one left: $x = 8, y = 2$. This is both a LPB and a WSB solution. The weak constraint has no impact on the solution in this case.

Now consider:

```
required  x + y = 10
strong    x ≥ 5
weak      y = 20
```

We'll only consider the WSB comparator this time, since it is more suitable for use with inequalities. (DeltaBlue does not handle inequalities.) Again, the required constraint has an infinite number of solutions. We winnow these down with the strong constraint to all $x \in [5, \infty)$, y such that $x + y = 10$. The weak constraint is unsatisfiable, so we minimize its error, resulting in the solution $x = 5, y = 5$.

Strict inequality constraints with metric comparators can be problematic in the presence of soft constraints, since they can lead to no solutions. Consider:

```
required  x > 10
weak      x = 5
```

This set of constraints has no solution — for any potential solution that satisfies $x > 10$, we can find another that better satisfies $x = 5$. For this reason, our examples usually use non-strict inequalities.

The “required” priority is special, in that those constraints must be satisfied in any solution. Both the semantics of hard and soft constraints, and the Cassowary and DeltaBlue solvers, can handle arbitrary numbers of soft priorities. However, for simplicity in the remainder of this note, we only use three, namely “strong,” “medium,” and “weak”.

We can also annotate variables used in constraints as *read-only*. Intuitively, when choosing the best solutions to a set of constraints with priorities, constraints should not be allowed to affect the choice of values for their read-only variables, i.e., information can flow out of the read-only variables, but not into them. Read-only annotations provide an important tool in a practical language for guiding the behavior of the constraint solver. However, they don’t present any particular issues for the formal semantics — we would just pass them on to the constraint solver to handle. Therefore, for simplicity we’ve dropped the extra rules to handle them (since they were just copies of the rules for variables without the read-only annotation).

3.1 Conjunctions and Disjunctions of Constraints

In general, a constraint might consist of conjunctions, disjunctions, and negations of atomic constraints. For a conjunction or disjunction, if there is a priority, it applies to the entire constraint, not to components. Thus this is legal:

strong $(x = 3 \vee x = 4)$

but this is not allowed:

$(\text{strong } x = 3) \vee (\text{weak } x = 4)$

Only some solvers, such as Z3, can accommodate disjunctions and explicit conjunctions of constraints. For DeltaBlue and Cassowary, conjunctions of constraints are implicitly specified by feeding multiple constraints to the solver, while disjunctions aren’t allowed. However, as noted above, the solver is a black box as far as the rest of the formal semantics is concerned.

3.2 Taming Identity Constraints

We introduce object identity in the Babelsberg/UID language, and continue to use it in our final language Babelsberg/Objects. A central issue in the design of a constraint object language is the interaction between constraints and object identity. Our earlier experience with Kaleidoscope suggests that it is all too easy to make constraints on object identity and types be too powerful, so that they lead to non-obvious consequences. This makes programs more difficult for the programmer to understand. To tame the power of constraints on object identity and type, we set the following goals for this aspect of our design:

1. In those languages that include object identity, we want to support explicit but straightforward identity constraints of the form $x==y$ (i.e., that variables x and y refer to the same object).
2. The solution to the constraints should be deterministic as far as object identity and type are concerned — there should never be multiple correct solutions in which a given variable refers to objects with different identities in the different solutions.
3. Any change to the identities of the objects referred to by variables should flow from an assignment statement — the constraint solver should not otherwise alter object identities.

4. To make Babelsberg programs more understandable for the programmer, we want to be able to reason about the constraints so that we can first solve all the identity constraints (with a deterministic solution, in keeping with Goal 3), and then the value constraints. A design that meets this goal of course benefits the language implementer as well.

In our intermediate language Babelsberg/Records, in which we introduce structured data in the form of immutable records but not yet object identity, we want analogous properties to hold: the solution to the constraints should be deterministic as far as object structure is concerned, and any changes to object structure should flow from an assignment statement.

4 Babelsberg/Reals and Babelsberg/PrimitiveTypes

We start with a very basic language, Babelsberg/Reals, that has only primitive values. In Babelsberg, constraints are expressions that return a boolean – the constraint solver’s task is to find values for the variables in the constraint’s expression such that it evaluates to true. So Boolean is a required datatype for all Babelsberg languages. In addition, we add reals as a second primitive type.

The evaluation model for Babelsberg/Reals is mostly standard for ordinary expressions and statements. In contrast, a statement that adds a constraint starts with a duration, namely **always** or **once**. The expression following the duration is taken unevaluated and added to the *constraint store*. For **always**, the constraint remains for the duration of the program’s execution; for **once**, it is removed after the solver finds a solution to the current set of constraints. The only wrinkle in the evaluation model for ordinary statements is that for an assignment statement, we evaluate the expression on the right hand side of the assignment, constrain the variable on the left hand side of the assignment to be equal to the result using a **once** constraint, and then turn all the constraints over to the solver. Doing this ensures that assignment interacts correctly with other constraints.

Here are some examples.

```
x := 3;  
x := 4;  
always x>=10
```

After evaluating the first statement we hand the following **once** constraint to the solver to find a value for x :

required $x = 3$

The solver finds a value for x , which is then used to update the environment to be $x \mapsto 3$.

Note that programs, as well as the variable names and values in the environment, are written in fixed pitch font. For contrast, we write the constraints that are handed to the solver in math font.

Continuing with the example, after the second statement we hand the following constraints to the solver:

weak $x = 3$
required $x = 4$

The weak $x = 3$ constraint is the stay constraint that x retain its previous value, while the required $x = 4$ constraint comes from the second assignment. This has the solution $x = 4$, resulting in a new environment $x \mapsto 4$.

The third statement adds the **always** constraint to the constraint store. So after that statement we have the following constraints:

```
weak   x = 4
required x ≥ 10
```

If we use a metric comparator such as weighted-sum-better or least-squares-better, we get the solution $x = 10$, since this minimizes the error for the weak constraint. If we use locally-predicate-better, then every $x \in [10, \infty)$ is a solution, and the system is free to select any of them. (However, as noted in Section 3, typically we wouldn't use LPB if we have inequalities.)

The following example illustrates using the same variable on the left and right hand sides of an assignment statement, as well as the interaction of assignments with `always` constraints.

```
x := 3;
y := 0;
always y = x+100;
x := x+2
```

After evaluating the first two statements and solving the resulting constraints, the environment has the binding $x \mapsto 3, y \mapsto 0$. The third statement causes the constraint `always y = x+100` to be added to the constraint store. We then hand the following constraints to the solver to find values for x and y :

```
weak   x = 3
weak   y = 0
required y = x + 100
```

This has multiple possible solutions, and the solver can select any one of them. Suppose it picks $x \mapsto 3, y \mapsto 103$.

After evaluating the next statement, we have the following constraints:

```
weak   x = 3
weak   y = 103
required y = x + 100
required x = 5
```

The first two constraints are the weak stays on x and y , the third comes from the `always` constraint in the constraint store, and the fourth comes from the assignment `x:=x+2` (where we evaluated `x+2` in the old environment to get 5). After solving these constraints, we have $x \mapsto 5, y \mapsto 105$.

New variables must be created with an assignment statement. Thus the following program is illegal — we would need to create x first with an assignment statement before adding the `always` constraint.

```
always x=10;
```

Requiring that variables be created before equating them with something can be annoying for the programmer. We decided that indeed the above program is illegal in the formal semantics. However, in practical implementations, we can have a shortcut to allow it. It only works for `=` constraints where one side is a new variable and all variables on the other side already exist. The behavior is that the system creates the new variable, evaluates the other side of the equality, assigns it to the new variable, and then adds the equality constraint as an `always` constraint.

An interesting aspect of this semantics, both as presented informally above and in the formalism that follows, is that we no longer model assignment as a constraint on variables at different times, as was done in for example the first Babelsberg paper [4]. See Appendix A.2 for a discussion of these alternatives. (This is

relegated to the appendix since, while interesting, it is largely orthogonal to the task at hand of providing a formal semantics for Babelsberg).

The program might also introduce simultaneous equations and inequalities. For example:

```
x := 0;
y := 0;
z := 0;
always x+y+2*z = 10;
always 2*x+y+z = 20;
x := 100
```

Assuming the solver can solve simultaneous linear equations, after the final assignment we will have $x \mapsto 100$, $y \mapsto -270$, $z \mapsto 90$.

As an example of unsatisfiable constraints, consider:

```
x := 5;
always x<=10;
x := x+15
```

After evaluating the first statement the environment includes the binding $x \mapsto 5$. After evaluating the statement that generates the `always` constraint, we solve the constraints

```
weak   x = 5
required x ≤ 10
```

This has the solution $x = 5$. Then we evaluate the last assignment, resulting the constraints

```
weak   x = 5
required x ≤ 10
required x = 20
```

Note that the required $x \leq 10$ constraint has persisted into this new set of constraints. These constraints are unsatisfiable.

4.1 Requirements for Constraint Expressions

The expressions that define constraints have a number of restrictions. These will apply to all of the Babelsberg languages.

1. Evaluating the expression that defines the constraint should return a boolean. (This is checked dynamically.)
2. The constraint expression should be free of side effects.¹
3. The result of evaluating the block should be deterministic. For example, an expression whose value depended on which of two processes happened to complete first wouldn't qualify. (This does not arise in the toy languages here, although we do need this restriction for a practical one.)

¹In a practical implementation, the programmer might be able to make cautious use of benign side effects in a constraint expression, for example, for caching or constructing temporary objects that are garbage collected before they are visible outside the constraint. In the formal semantics, however, we simply disallow side effects in constraint expressions.

4.2 Control Structures

Babelsberg/Reals includes `if` and `while` control structures. These work in the usual way, and allow (for example) a variable to be incremented only if a test is satisfied, or an `always` constraint to be conditionally asserted. The test for an `if` statement is evaluated, and one or the other branch is taken — there is no notion of backtracking to try the other branch. (Adding Prolog-style backtracking is left for future work — see Appendix A.6.) Similarly, a `while` statement executes the body a fixed number of times — there is no possibility of backtracking to execute it a different number of times.

The test in an `if` or `while` statement can use short-circuit evaluation when evaluating an expression involving `and` and `or`. For example, this program results in $x \mapsto 100$ (and doesn't get a divide-by-zero error):

```
x := 4;
if x=4 or x/0=10
  then x := 100
  else x := 200
```

For simplicity, our formal rules don't include short-circuit evaluation — adding it would be straightforward but would require additional, not-very-interesting rules.

Constraints with conjunctions or disjunctions are just turned over to the solver, rather than being evaluated using short-circuit evaluation. We could also add `if` expressions to the language (distinct from `if` statements). However, since there is a simple translation from `if` expressions to conjunctions and disjunctions, we don't include them. (If we did have them, they would also need to simply be turned over to the solver.) For example, the following two constraints are equivalent, and have the solution $x \mapsto 10$:

```
always if x=4 then x=5 else x=10

always (x=4 and x=5) or (x!=4 and x=10)
```

In either case, we would turn the entire constraint over to the solver to find a solution for x .

Here is an example of an unsatisfiable constraint of this sort:

```
always if x=4 then x=5 else x=4
```

which is equivalent to:

```
always (x=4 and x=5) or (x!=4 and x=4)
```

4.3 Adding Other Primitive Types

It is straightforward to extend Babelsberg/Reals with other primitive types, such as integers and strings. When we need to refer to this language rather than just Babelsberg/Reals we will call it Babelsberg/PrimitiveTypes. Note that all the types in Babelsberg/PrimitiveTypes are atomic — we don't have recursive types or types that define values that hold other values (such as records, arrays, or sets).

Since we are modeling a dynamically typed language, we need to consider the case of changing the type of a variable. Here's a simple example program that illustrates this.


```
x := 5;
x := "Hello";
```

After the second statement we solve the following constraints:

```
weak   x = 5
required x = "Hello"
```

The final result is $x \mapsto \text{"Hello"}$.

The changed type might propagate through an `always` constraint:

```
x := 5;
y := 10;
always y = x;
x := "Hello";
```

The final result is that both `x` and `y` are strings.

There is also a potential interaction with overloaded operators. Suppose that for strings `+` denotes string concatenation:

```
x := 5;
y := 10;
always y = x+x;
x := "Hello";
```

After the `always` statement we have $x \mapsto 5$, $y \mapsto 10$. Then after the final statement, we have the constraints

```
weak   x = 5
weak   y = 10
required y = x + x
required x = "Hello"
```

This has the solution $x \mapsto \text{"Hello"}$, $y \mapsto \text{"HelloHello"}$.

Non-determinism can arise for primitive types as well as values:

```
x := 3;
always weak x=5;
always weak x="hello";
```

Here the solver is free to choose either an integer or a string for `x`. This is strange, and also not really in keeping with the spirit of the goals for our languages with respect to constraints and object identity (Section 3.2). However, it is just an artifact of the semantic rules we are using for Babelsberg/PrimitiveTypes — for simplicity, in these rules we don't distinguish among primitive types, with the consequence that the above program is legal. When we get to our final goal (the Babelsberg/Objects language) the corresponding program will be disallowed due to the automatic boxing and unboxing of primitive types in the formal semantic rules.

4.4 Formalism

We present the formal semantics of Babelsberg/PrimitiveTypes.

4.4.1 Syntax

Statement	s	$::=$	<code>skip</code> <code>x := e</code> <code>always C</code> <code>once C</code> <code>s; s</code> <code>if e then s else s</code> <code>while e do s</code>
Constraint	C	$::=$	ρe $C \wedge C$
Expression	e	$::=$	c x $e \oplus e$ $e \otimes e$ $e \odot e$
Constant	c	$::=$	<code>true</code> <code>false</code> base type constants
Variable	x	$::=$	variable names
Value	v	$::=$	c

The language includes a set of boolean and base type constants (e.g., reals), ranged over by metavariable c . A finite set of operators on expressions is ranged over by \oplus . For booleans this set includes no operations, but it does include operations on the reals such as $+$ and $*$. The symbol \otimes ranges over a set of *predicate* operators ($=$ and \neq for booleans, \leq , $<$, $=$, and so on for reals.) These test properties of base types in the language. The symbol \odot ranges over a set of logical operators for combining boolean expressions (e.g., \wedge , \vee). The predicate operators are assumed to include at least an equality operator $=$ for each primitive type in the language, and the logical operators are assumed to include at least conjunction \wedge . The syntax of this language does have some limitations as compared with that of a practical language — for example, there are only binary operators (not unary or ternary), and the result must have the same type as the arguments. We make these simplifications since the purpose of Babelsberg/PrimitiveTypes is to elucidate the semantics of such languages as a step toward Babelsberg/Objects, rather than to specify a real language.

For constraints, the symbol ρ ranges over a finite and totally ordered set of constraint *priorities* and is assumed to include a bottom element *weak* and a top element *required*. While syntax requires the priority to be explicit, for simplicity we sometimes omit it in this semantics. A constraint with no explicit priority implicitly has the priority *required*. Finally, for simplicity we do not model read-only annotations in the formal semantics.

The syntax is thus that of a simple, standard imperative language except for the **always** and **once** statements, which declare constraints. An **always** constraint must hold for the rest of the programs execution, whereas a **once** constraint is satisfied by the solver and then retracted. Note that for simplicity this semantics implicitly gets stuck whenever the solver cannot satisfy a constraint, either due to an unsatisfiable constraint or due to the solver being unable to determine whether the constraint is satisfiable. In a practical implementation, we would likely want to differentiate between these cases, since it's useful if we can inform the programmer that the constraints are truly not satisfiable. We could also add standard exception handling to remove the unsatisfiable or unknown constraint and continue, but omit this here for simplicity.

4.4.2 Semantics

The semantics is defined by several judgments, defined below. These judgments depend on the notion of an *environment*, which is a partial function from program variables to program values. Metavariable E ranges over environments. When convenient we also view an environment as a set of (program variable, program value) pairs. For each operator o in the language we assume the existence of a corresponding semantic function denoted $\llbracket o \rrbracket$.

$$\boxed{E \vdash e \Downarrow v}$$

“Expression e evaluates to value v in the context of environment E .”

$$E \vdash c \Downarrow c \qquad \text{(E-CONST)}$$

$$\begin{array}{c}
\frac{\mathbf{E}(x) = v}{\mathbf{E} \vdash x \Downarrow v} \quad (\text{E-VAR}) \\
\\
\frac{\mathbf{E} \vdash e_1 \Downarrow v_1 \quad \mathbf{E} \vdash e_2 \Downarrow v_2 \quad v_1 \llbracket \oplus \rrbracket v_2 = v}{\mathbf{E} \vdash e_1 \oplus e_2 \Downarrow v} \quad (\text{E-OP}) \\
\\
\frac{\mathbf{E} \vdash e_1 \Downarrow v_1 \quad \mathbf{E} \vdash e_2 \Downarrow v_2 \quad v_1 \llbracket \ominus \rrbracket v_2 = v}{\mathbf{E} \vdash e_1 \ominus e_2 \Downarrow v} \quad (\text{E-COMPARE}) \\
\\
\frac{\mathbf{E} \vdash e_1 \Downarrow v_1 \quad \mathbf{E} \vdash e_2 \Downarrow v_2 \quad v_1 \llbracket \odot \rrbracket v_2 = v}{\mathbf{E} \vdash e_1 \odot e_2 \Downarrow v} \quad (\text{E-COMBINE})
\end{array}$$

$$\boxed{\mathbf{E} \models \mathbf{C}}$$

This judgment represents a call to the constraint solver, which we treat as a black box. The proposition $\mathbf{E} \models \mathbf{C}$ denotes that environment \mathbf{E} is a *solution* to the constraint \mathbf{C} (and further one that is optimal according to the solver’s semantics, as discussed earlier).

$$\boxed{\text{stay}(\mathbf{E}) = \mathbf{C}}$$

This judgment defines how to translate an environment into a source-level “stay” constraint.

$$\begin{array}{c}
\text{stay}(\emptyset) = \text{true} \quad (\text{STAYEMPTY}) \\
\\
\frac{\mathbf{E}(x) = v \quad \mathbf{E}_0 = \mathbf{E} \setminus \{(x, v)\} \quad \text{stay}(\mathbf{E}_0) = \mathbf{C}_0 \quad \mathbf{C} = \mathbf{C}_0 \wedge \text{weak } x=v}{\text{stay}(\mathbf{E}) = \mathbf{C}} \quad (\text{STAYONE})
\end{array}$$

$$\boxed{\langle \mathbf{E} | \mathbf{C} | s \rangle \longrightarrow \langle \mathbf{E}' | \mathbf{C}' \rangle}$$

“Execution starting from configuration $\langle \mathbf{E} | \mathbf{C} | s \rangle$ ends in state $\langle \mathbf{E}' | \mathbf{C}' \rangle$.”

A “configuration” defining the state of an execution includes a concrete context, represented by the environment, a symbolic context, represented by the constraint, and a statement to be executed. The environment and statement are standard, while the constraint is not part of the state of a computation in most languages. Intuitively, the environment comes from constraint solving during the evaluation of the immediately preceding statement, and the constraint records the **always** constraints that have been declared so far during execution. Note that our execution implicitly gets stuck if the solver cannot produce a model.

$$\frac{\mathbf{E} \vdash e \Downarrow v \quad \text{stay}(\mathbf{E}) = \mathbf{C}_s \quad \mathbf{E}' \models (\mathbf{C} \wedge \mathbf{C}_s \wedge x = v)}{\langle \mathbf{E} | \mathbf{C} | x := e \rangle \longrightarrow \langle \mathbf{E}' | \mathbf{C} \rangle} \quad (\text{S-ASGN})$$

$$\frac{\text{stay}(\mathbf{E}) = \mathbf{C}_s \quad \mathbf{E}' \models (\mathbf{C} \wedge \mathbf{C}_s \wedge \mathbf{C}_0)}{\langle \mathbf{E} | \mathbf{C} | \text{once } \mathbf{C}_0 \rangle \longrightarrow \langle \mathbf{E}' | \mathbf{C} \rangle} \quad (\text{S-ONCE})$$

$$\frac{\langle \mathbf{E} | \mathbf{C} | \text{once } \mathbf{C}_0 \rangle \longrightarrow \langle \mathbf{E}' | \mathbf{C} \rangle \quad \mathbf{C}' = \mathbf{C} \wedge \mathbf{C}_0}{\langle \mathbf{E} | \mathbf{C} | \text{always } \mathbf{C}_0 \rangle \longrightarrow \langle \mathbf{E}' | \mathbf{C}' \rangle} \quad (\text{S-ALWAYS})$$

$$\begin{array}{c}
\langle E|C|\text{skip}\rangle \longrightarrow \langle E|C\rangle \quad (\text{S-SKIP}) \\
\\
\frac{\langle E|C|s_1\rangle \longrightarrow \langle E'|C'\rangle \quad \langle E'|C'|s_2\rangle \longrightarrow \langle E''|C''\rangle}{\langle E|C|s_1;s_2\rangle \longrightarrow \langle E''|C''\rangle} \quad (\text{S-SEQ}) \\
\\
\frac{E \vdash e \Downarrow \text{true} \quad \langle E|C|s_1\rangle \longrightarrow \langle E'|C'\rangle}{\langle E|C|\text{if } e \text{ then } s_1 \text{ else } s_2\rangle \longrightarrow \langle E'|C'\rangle} \quad (\text{S-IFTHEN}) \\
\\
\frac{E \vdash e \Downarrow \text{false} \quad \langle E|C|s_2\rangle \longrightarrow \langle E'|C'\rangle}{\langle E|C|\text{if } e \text{ then } s_1 \text{ else } s_2\rangle \longrightarrow \langle E'|C'\rangle} \quad (\text{S-IFELSE}) \\
\\
\frac{E \vdash e \Downarrow \text{true} \quad \langle E|C|s\rangle \longrightarrow \langle E'|C'\rangle \quad \langle E'|C'|\text{while } e \text{ do } s\rangle \longrightarrow \langle E''|C''\rangle}{\langle E|C|\text{while } e \text{ do } s\rangle \longrightarrow \langle E''|C''\rangle} \quad (\text{S-WHILED O}) \\
\\
\frac{E \vdash e \Downarrow \text{false}}{\langle E|C|\text{while } e \text{ do } s\rangle \longrightarrow \langle E|C\rangle} \quad (\text{S-WHILESKIP})
\end{array}$$

5 Babelsberg/Records

For this next language, we augment Babelsberg/PrimitiveTypes with immutable records. This language is interesting on its own as an expository language and as a step toward Babelsberg/Objects; but we will also continue to use these immutable records in the formal semantics for Babelsberg/Objects to represent *value classes* (see Section 7).

Records are written as lists of name/value pairs in curly braces:

```
{x:5, y:10}
```

There is no notion of object identity for Babelsberg/Records — we can test whether two records are equal, but whether or not they are identical would be an implementation issue and not part of the semantics.

The syntax is extended to include record constructors and field access. Here are examples of expressions, assignment statements, and constraints involving records:

```
p := {x:2, y:5}; /* assign a record to p */
a := p.x;      /* access a field of a record and assign it to a variable */
q := p;        /* q is now a copy of p (or maybe it's shared; we can't tell
               the difference) */
always p.x = 100; /* a constraint on a record field */
always q = p;    /* an equality constraint between two records */
always q.y = 20;
```

Note the difference between the assignment `q:=p` and the constraint `always q=p`. After the `always p.x=100` constraint, we have `p ↦ {x:100, y:5}` but `q ↦ {x:2, y:5}` — at this point `p` and `q` are unrelated, so adding the constraint on `p.x` had no effect on `q`. However, after the final constraint, we have `p ↦ {x:100, y:20}` and `q ↦ {x:100, y:20}`, since `p` and `q` are now constrained to be equal.

The solver must now also handle records. To tame the power of the solver so that it does not (for example) invent new fields for records, we add *structural compatibility checks* on constraints. These structural

compatibility checks are assertions that are checked dynamically before sending the constraints involving records to the solver, for example, checking whether a variable is bound to a record, and whether the record has the necessary fields. While these assertions are checked, unlike constraints the system will never change anything to enforce them — if one is violated it's just an error. Instead, the programmer must ensure that a record with the expected fields is first assigned to a variable used in record constraints, just as a programmer would need to ensure that a record with the expected fields was assigned to a record-valued variable in a standard language.

Here are a few examples of structural compatibility checks.

```
p := {x:2, y:5};
always p.x = 100;
```

The structural compatibility check is that `p` is a record that has an `x` field, which succeeds.

The following program is OK — just as in `Babelsberg/PrimitiveTypes`, we can change the type of a variable using an assignment.

```
a := {x:1};
a := {y:10};
```

However, in contrast to `Babelsberg/PrimitiveTypes`, the following program fails the structural compatibility check — only an assignment can change the type of a variable.

```
a := {x:1};
once a = {y:10};
```

This program fails as well:

```
a := {x:1};
b := {x:1};
always a=b;
a := {x:1, y:10};
```

Here, assigning a record with a different structure to `a` is OK on its own, but the `always` constraint would also require `b` to change. It's a bit weird that this behavior is different from the behavior with primitive types (in which a similar program was OK). However, once we introduce object identity, we will be able to have changes to the types of variables ripple through the system via identity constraints (but it will need to be via identity constraints and not value constraints).

A few more examples:

These two programs fail, because as in `Babelsberg/PrimitiveTypes` only assignment can initialize variables.

```
once a = {y:10};
```

```
a := {y:10};
always b=a;
```

The following program also fails the structural compatibility check, since the `always` constraint expects `p` to have a `y` field but it doesn't:

```
p := {x:2};
always p.y = 100;
```

This program fails as well, since constraining a record to be equal to a number fails the compatibility check:

```
p := {x:2};
always p = 5;
```

The following program passes the structural compatibility checks, but fails with an unsatisfiable constraint error, since we are requiring that `p.x` be both 100 and 2:

```
p := {x:0, y:0};
always p.x = 100;
p := {x:2, y:5};
```

However, the following program is OK:

```
p := {x:0};
always medium p = {x:3};
always medium p = {x:4}
```

The solution will be that `p` is a record with a single field `x` bound to a number, with the exact value depending on the solver (3, 4, or 3.5 being the most reasonable possibilities). Note that the solver may need to adjudicate among solutions that have different primitive types in a given field, but the structural compatibility checks ensure that it doesn't need to decide between two records with e.g. different numbers of fields.

Here's another example that fails a structural compatibility check:

```
a := {x:0};
b := {y:5};
always a=b
```

This is the case even though there actually are records that would satisfy the required `a=b` constraint — the issue is that the assignment to `a` leaves it with a record type that has a single `x` field, the assignment to `b` leaves it with a record type that has a single `y` field, and the structural compatibility check prevents the solver from changing either of these types. This illustrates one aspect of using the structural compatibility checks to tame the solver — one could otherwise imagine the solver coming up with the solution `a ↦ {x:0, y:5}` and `b ↦ {x:0, y:5}`, or even `a ↦ {x:0, y:5, z:10}` and `b ↦ {x:0, y:5, z:10}`. See Appendix A.4 for more about this issue.

5.1 Formalism

5.1.1 Syntax

The syntax from Section 4 is augmented now to support records and the ability to access fields of a record:

Expression $e ::= \dots \mid \{l_1:e_1, \dots, l_n:e_n\} \mid e.l$
 Label $l ::=$ record label names

Value $v ::= c \mid \{l_1:v_1, \dots, l_n:v_n\}$

We now assume the solver “understands” records and record operations directly. We assume the equality operator $=$ can be used to compare two records for logical equality, but no other operators apply to record values.

5.1.2 Semantics

$E \vdash e \Downarrow v$

Expression evaluation is updated to support records.

$$\frac{E \vdash e_1 \Downarrow v_1 \cdots E \vdash e_n \Downarrow v_n}{E \vdash \{l_1:e_1, \dots, l_n:e_n\} \Downarrow \{l_1:v_1, \dots, l_n:v_n\}} \quad (\text{E-REC})$$

$$\frac{E \vdash e \Downarrow \{l_1:v_1, \dots, l_n:v_n\} \quad 1 \leq i \leq n}{E \vdash e.l_i \Downarrow v_i} \quad (\text{E-FIELD})$$

$E \vdash e : T$

$E \vdash C$

“Expression e has type T in the context of environment E .”

“Constraint C is well formed in the context of environment E .”

We use a notion of typechecking to prevent undesirable non-determinism in constraints. Specifically, we want constraint solving to preserve the structure of the values of variables, changing only the underlying primitive data as part of a solution, in support of the goals listed in Section 3.2. We formalize our notion of structure through a simple syntax of types:

Type $T ::= \text{PrimitiveType} \mid \{l_1:T_1, \dots, l_n:T_n\}$

The typechecking rules are mostly standard. However, we check expressions dynamically just before constraint solving, so we typecheck in the context of a runtime environment, which is somewhat unusual.

$$E \vdash c : \text{PrimitiveType} \quad (\text{T-CONSTANT})$$

$$\frac{E \vdash x \Downarrow v \quad E \vdash v : T}{E \vdash x : T} \quad (\text{T-VARIABLE})$$

$$\frac{E \vdash e_1 : \text{PrimitiveType} \quad E \vdash e_2 : \text{PrimitiveType}}{E \vdash e_1 \oplus e_2 : \text{PrimitiveType}} \quad (\text{T-OP})$$

$$\frac{E \vdash e_1 : T \quad E \vdash e_2 : T}{E \vdash e_1 \otimes e_2 : \text{PrimitiveType}} \quad (\text{T-COMPARE})$$

The above rule ensures that two records can only be compared for equality when they have the same structure. This is necessary to ensure that the solver will never have to invent record fields in order to satisfy a record equality constraint.

$$\frac{E \vdash e_1 : \text{PrimitiveType} \quad E \vdash e_2 : \text{PrimitiveType}}{E \vdash e_1 \odot e_2 : \text{PrimitiveType}} \quad (\text{T-COMBINE})$$

$$\frac{E \vdash e_1 : T_1 \cdots E \vdash e_n : T_n}{E \vdash \{l_1 : e_1, \dots, l_n : e_n\} : \{l_1 : T_1, \dots, l_n : T_n\}} \quad (\text{T-REC})$$

$$\frac{E \vdash e : \{l_1 : T_1, \dots, l_n : T_n\} \quad 1 \leq i \leq n}{E \vdash e.l_i : T_i} \quad (\text{T-FIELD})$$

The above rule ensures that a field of a record can only be referenced when it already exists in the record. This is necessary to ensure that the solver will never have to invent record fields in order to satisfy a constraint containing field accesses.

The following two rules simply ensure that a constraint is well typed.

$$\frac{E \vdash e : T}{E \vdash \rho e} \quad (\text{T-PRIORITY})$$

$$\frac{E \vdash C_1 \quad E \vdash C_2}{E \vdash C_1 \wedge C_2} \quad (\text{T-CONJUNCTION})$$

$$\boxed{\langle E | C | s \rangle \longrightarrow \langle E' | C' \rangle}$$

The semantics for executing statements is essentially identical to what we had before, except that now we typecheck constraints before we solve them. Implicitly we get “stuck” if such a check fails. In a practical language, this could be extended to generate an exception instead. We only show the modified rules below.

$$\frac{E \vdash e \Downarrow v \quad E[x \mapsto v] \vdash C \quad \text{stay}(E) = C_s \quad E' \models (C \wedge C_s \wedge x = v)}{\langle E | C | x := e \rangle \longrightarrow \langle E' | C \rangle} \quad (\text{S-ASGN})$$

The second premise above is necessary so that we don’t pass the solver an ill-typed set of constraints. Specifically, since the assignment can change the type of x , we have to make sure that the global constraints in scope, C , are well typed with respect to this new type. Note that the stay constraint for x within C_s may not be well typed with respect to the new type of x , but that’s OK since the required constraint $x = v$ will take precedence. The syntax $E[x \mapsto v]$ denotes the environment identical to E but with x mapped to the value v .

$$\frac{E \vdash C_0 \quad \text{stay}(E) = C_s \quad E' \models (C \wedge C_s \wedge C_0)}{\langle E | C | \text{once } C_0 \rangle \longrightarrow \langle E' | C \rangle} \quad (\text{S-ONCE})$$

$$\frac{\langle E | C | \text{once } C_0 \rangle \longrightarrow \langle E' | C \rangle \quad C' = C \wedge C_0}{\langle E | C | \text{always } C_0 \rangle \longrightarrow \langle E' | C' \rangle} \quad (\text{S-ALWAYS})$$

5.2 Adding Mutable Records

It would be easy to extend Babelsberg/Records to allow mutable records. Syntactically, we would simply allow field accesses as l-values, e.g.,

```
p := {x:0, y:0};  
p.x := 100;
```

After the second assignment, we would have $p \mapsto \{x:100, y:0\}$. This doesn't add any particular complications to the semantics.

However, the only additional feature provided by such an extension would be the syntax allowing field accesses as l-values — we can always convert such a program into one that only has immutable records. For example, the above program is equivalent to the following program in Babelsberg/Records:

```
p := {x:0, y:0};  
once p.x = 100;
```

In any case there still would be no notion of object identity. Consider:

```
p := {x:0, y:0};  
q := p;  
p.x := 100;
```

After executing the three statements $p \mapsto \{x:100, y:0\}$, but $q \mapsto \{x:0, y:0\}$.

Mutable records would thus simply provide a syntactic convenience, rather than some new capability. We would still want immutable records as well (which we will use to model instances of value classes in the formal semantics for Babelsberg/Objects), so this would also add clutter to the languages. We therefore don't include Babelsberg with mutable records as a separate language.

6 Babelsberg/UID

Babelsberg/UID adds a number of features to the language. As another step toward representing objects, we augment Babelsberg/PrimitiveTypes with records that live on the heap, and that are mutable, have an identity, and can be aliased. As with Babelsberg/Records these are represented as lists of name/value pairs in curly braces:

```
{x:5, y:10}
```

However, to emphasize that we now allocate records with object identity on the heap, we use the keyword `new` when creating one. The syntax includes field accesses as both r-values and l-values.

In all the examples in this section, we will only use records with object identity. However, later, when we are formalizing Babelsberg/Objects, we will use both kinds of records: records with object identity to model objects that have identity, and immutable records (as in Babelsberg/Records) to model instances of value classes (which are immutable objects for which object identity is not significant). It should be clear from context when we use the word “record” what kind we mean; but when it's necessary to distinguish

them, we'll use the terms "i-record" and "uid-record" for immutable records and records with object identity respectively.

Only uid-records live on the heap — primitive values and i-records do not. (It would be possible to store all data on the heap, but we elected not to, since having everything on the heap makes the descriptions more complex — there would always be a level of indirection to get to data. If one wants the effect of storing an integer or boolean or i-record on the heap, it is easy to simulate this by constructing a uid-record that has a single field that holds the integer or boolean or i-record.)

A field of a uid-record holds either a primitive type, a i-record, or a reference to another uid-record. There is no syntax for creating a nested uid-record directly — nested records have to be constructed with references, so each record *must* have a variable that refers to it, even if it is only used in a nested structure. This simplifies the semantics. (Nested uid-records could be supported directly with just a source code transformation. In any case, when we get to Babelsberg/Objects, we will remove this restriction and allow expressions that do involve creating new objects.)

Since object identity is now significant, we add identity constraints to the language (following Smalltalk syntax, written `==`, in contrast to `=` for equality constraints). For records `p` and `q`, if `p` and `q` are identical, they must also be equal, but the converse is not necessarily true — if `p` and `q` are equal, they might or might not be identical.

Unlike Babelsberg/PrimitiveTypes and Babelsberg/Records, the weak stays on variables are identity constraints rather than equality constraints. For uid-records, such a variable continues to refer to the same object unless reassigned. This is a direct consequence of the weak stays referring to the references the variables hold, not the records on the heap. In addition, there are also weak stay constraints on the values in the record fields.

For simplicity, we also allow identity tests and constraints on primitive types and i-records. (Otherwise we would need two different translations for assignment.) Again for simplicity, two instances of a primitive type or i-record are identical iff they are equal. (Allowing them to be equal but not identical would seem to imply storing them on the heap, which would complicate the semantics.)

Here are some examples involving uid-record fields and constraints:

```
p := new {x:2, y:5}; /* create a record on the heap and save the reference in p */
a := p.x;          /* read a field of a record and assign it to a variable */
p.x := 6;          /* assign 6 to the field p.x */
always p.x = 100;
```

As with Babelsberg/Records, there is no static type checking — checking is all done dynamically. As before, our model includes structural compatibility checks on constraints, which tames the solver so that it doesn't do such things as inventing new records or adding fields to a record. (The only way to create a new record is with an explicit `new` expression in the program, and this also defines all of the fields that it has and will ever have.)

After executing the right hand side of the first statement, the heap is updated with a new object and a reference `r` to point to it. The assignment and stay constraints then desugar into:

```
required  p = r
required  H(r) = {x : xr, y : yr}
weak     xr = 2
weak     yr = 5
```

The reference `r` is a constant. We assume the solver understands uninterpreted function symbols. The uninterpreted function symbol `H` in the second constraint above is used to represent the heap's mapping

from references to their contents; the constraint forces the solver to keep the reference r pointing to a record that has *only* x and y fields. We introduce variables for the values of these fields and add weak stays so these keep the values with which they were created – these can always be satisfied, because the variables are fresh.

After the second statement we have the stay constraints on p and its fields, and the required constraint resulting from the assignment:

```

weak   p = r
required H(r) = {x : x_r, y : y_r}
weak   x_r = 2
weak   y_r = 5
required a = 5

```

After the third statement:

```

weak   p = r
required H(r) = {x : x_r, y : y_r}
weak   x_r = 2
weak   y_r = 5
weak   a = 5
required p.x = 6

```

And finally:

```

weak   p = r
required H(r) = {x : x_r, y : y_r}
weak   x_r = 6
weak   y_r = 5
weak   a = 5
required p.x = 100

```

After solving the final set of constraints we have the environment $p \mapsto r \wedge a \mapsto 5$, as well as the heap $r \mapsto \{x:100, y:5\}$.

Here is an example that fails the structural compatibility checks.

```

p := new {x:2, y:5};
always p.z = 5;

```

As before, the system is not allowed to add fields to records, so p would be required to have a z field already.

The following example demonstrates one form of aliasing:

```

p := new {x:2, y:5};
q := p;
p.x := 100;
q := new {z:10};
p.x := 200;

```

After the assignment $q:=p$, q holds the same reference as p , so the subsequent assignment to $p.x$ changes both. But then we break the alias with another assignment to q , so the second assignment to $p.x$ doesn't affect q (which by that point no longer has an x field).

There can also be explicit identity constraints:

```
p := new {x:2, y:5};
q := p;
always q==p;
q := new {z:10};
```

After the last statement both `p` and `q` point to the same `{z:10}` record. Note that we don't have structural compatibility checks for re-satisfying identity constraints, so that the assignment to `q` that changes its structure ripples through the identity constraint to change `p` as well. However, if the `always` constraint had been a record equality constraint rather than an identity constraint, the last assignment to `q` would have failed the structural compatibility check, which would have expected it to be a record with `x` and `y` fields.

In Section 3.2 we listed a number of goals for our design to attempt to tame the power of identity constraints and the solver. In support of this, we require that variables be created in an assignment statement — the above program would have been illegal if we didn't have the `q:=p` statement but just tried to create `q` with the `always` constraint. We also require that a new identity constraint be satisfied at the time it is created — otherwise it's an error (modeled in the formal semantics by the rules getting stuck). This restriction was obeyed in the above program. Here is an example that violates it, and is hence not legal:

```
p := new {x:2};
q := new {y:5}
always q==p;
```

Without the restriction, after the last statement `p` and `q` would definitely point to the same record, but it might be either `{x:2}` or `{y:5}`. (In any case, `p` and `q` would still one of the records that was already created by the `new` expressions — the solver wouldn't invent a new one — but we want to tame the language further.)

With the restriction, any changes to the types of variables must flow from an explicit assignment statement — the effects can still ripple outwards via identity constraints — but they are deterministic. For the same reason, we disallow disjunctions of identity constraints (since otherwise there could be programs with multiple correct solutions to the identity constraints). It is only updates to values that can be non-deterministic as the result of constraints. One item for future work will be to formalize and prove this statement, which seems like a useful property for the language.

As another aspect of taming identity constraints, explicit priorities on identity constraints are not allowed — there are only the implicit weak identity stay constraints on all variables. (We haven't found any compelling use cases for programmer-specified soft identity constraints, and omitting them simplifies reasoning about identity constraints for the programmer.) Similarly, we disallow disjunctions of identity constraints, which would also introduce nondeterminism; and for simplicity, there are no explicit conjunctions of identity constraints, only the implicit conjunction resulting from writing several of them.

This program is thus illegal:

```
p := new {x:0};
q := new {x:5};
always medium p.x = 0;
always medium.x = 5;
always weak p==q
```

Without the restriction, the programmer would need to interleave determining values and identities in reasoning about the program's behavior. For example, for the above program, to decide that we should leave

the weak identity constraint unsatisfied, we'd need to first solve the value constraints and then decide that we can't satisfy the weak identity constraint.

The solver can't spontaneously create new uid-records if they are needed — new uid-records can only be created with an explicit assignment with a `new` on the right hand side. Consider:

```
a := new {x:1};
b := a;
always a.x=1;
always b.x=2;
```

After `b:=a`, `a` and `b` refer to the same record, but after the second `always` constraint, the solver would need to spontaneously create a new record and point `b` at it to satisfy the constraints. So this program fails.

It also can't switch identities around to satisfy value constraints. Here is the same program, except that there happens to be a record with the required value for its `x` field already lying around. But to no avail: this program fails as well, rather than silently changing `b` to refer to the `{x:2}` record.

```
a := new {x:1};
b := a;
c := new {x:2};
always a.x=1;
always b.x=2;
```

6.1 Formalism

Since this version is significantly different than the prior formalisms, we present it in its entirety rather than as a delta from those ones.

6.1.1 Syntax

The syntax is augmented to create records on the heap with the `new` keyword. We also add references as values and an identity operation `==` to compare values by identity rather than by structure (for primitive values, this is the same as `=`). The arguments to identity operations are not arbitrary expressions but rather *accessors*, which are essentially access paths through a reference. This is no loss of expressiveness but is useful for the translation to the constraint solver.

Statement	<code>s</code>	<code>::= skip a := e x := new o always C once C s;s</code> <code> if e then s else s while e do s</code>
Constraint	<code>C</code>	<code>::= ρ e C ∧ C</code>
Expression	<code>e</code>	<code>::= c r a e ⊕ e e ⊗ e e ⊙ e I D</code>
Identity	<code>I</code>	<code>::= a == a</code>
Accessor	<code>a</code>	<code>::= x a.l</code>
Record	<code>o</code>	<code>::= {l₁:e₁, ..., l_n:e_n}</code>
Dereference	<code>D</code>	<code>::= H(r)</code>
Reference	<code>r</code>	<code>::= references to heap records</code>
Constant	<code>c</code>	<code>::= true false base type constants</code>
Variable	<code>x</code>	<code>::= variable names</code>
Value	<code>v</code>	<code>::= c r</code>

In the syntax, we treat H as a keyword. Source programs will not use expressions of the form $H(\mathbf{r})$, but they are introduced as part of constraints given to the solver, which we assume will treat H as an uninterpreted function, as described earlier.

6.1.2 Semantics

In addition to an environment, the semantics now requires a *heap*, which is a partial function from mutable references to record “objects” of the form $\{l_1:v_1, \dots, l_n:v_n\}$. Heaps are ranged over by metavariable H^2 . Secondly, in addition to the normal active constraints, we also now have a number of required identity constraints that are active. These are represented by I .

$$\boxed{E; H \vdash e \Downarrow v}$$

“Expression e evaluates to value v in the context of environment E and heap H .”

Most of the updates to the expression rules just use the new judgment that includes the heap. We do not include rules for evaluating expressions of the form $H(\mathbf{r})$, since they are not intended to appear in source programs.

$$E; H \vdash c \Downarrow c \quad (\text{E-CONST})$$

$$\frac{E(x) = v}{E; H \vdash x \Downarrow v} \quad (\text{E-VAR})$$

$$\frac{E; H \vdash a \Downarrow r \quad H(\mathbf{r}) = \{l_1:v_1, \dots, l_n:v_n\} \quad 1 \leq i \leq n}{E; H \vdash a.l_i \Downarrow v_i} \quad (\text{E-FIELD})$$

$$E; H \vdash r \Downarrow r \quad (\text{E-REF})$$

$$\frac{E; H \vdash e_1 \Downarrow v_1 \quad E; H \vdash e_2 \Downarrow v_2 \quad v_1 \llbracket \oplus \rrbracket v_2 = v}{E; H \vdash e_1 \oplus e_2 \Downarrow v} \quad (\text{E-OP})$$

$$\frac{E; H \vdash e_1 \Downarrow v_1 \quad E; H \vdash e_2 \Downarrow v_2 \quad v_1 \llbracket \ominus \rrbracket v_2 = v}{E; H \vdash e_1 \ominus e_2 \Downarrow v} \quad (\text{E-COMPARE})$$

$$\frac{E; H \vdash e_1 \Downarrow v_1 \quad E; H \vdash e_2 \Downarrow v_2 \quad v_1 \llbracket \odot \rrbracket v_2 = v}{E; H \vdash e_1 \odot e_2 \Downarrow v} \quad (\text{E-COMBINE})$$

$$\frac{E; H \vdash a_1 \Downarrow v \quad E; H \vdash a_2 \Downarrow v}{E; H \vdash a_1 == a_2 \Downarrow \text{true}} \quad (\text{E-IDENTITYTRUE})$$

$$\frac{E; H \vdash a_1 \Downarrow v_1 \quad E; H \vdash a_2 \Downarrow v_2 \quad v_1 \neq v_2}{E; H \vdash a_1 == a_2 \Downarrow \text{false}} \quad (\text{E-IDENTITYFALSE})$$

$$\boxed{E; H \vdash e : T}$$

²Not to be confused with the keyword H in the syntax; it is always clear from context which is being used.

$$\boxed{E;H \vdash C}$$

The typing rules also are augmented with the heap. They are otherwise unchanged. The additional thing to type are references. We do not add type rules for identities. This ensures that constraints involving them do not typecheck, so identity checks cannot occur in ordinary constraints. As with the evaluation rules, expressions of the form $H(\mathbf{r})$ do not typecheck.

$$\begin{array}{c} E;H \vdash c : \text{PrimitiveType} \quad (\text{T-CONSTANT}) \\ \\ \frac{E;H \vdash o : \{l_1:T_1, \dots, l_n:T_n\} \quad H(\mathbf{r})=o}{E;H \vdash \mathbf{r} : \{l_1:T_1, \dots, l_n:T_n\}} \quad (\text{T-REF}) \\ \\ \frac{E \vdash x \Downarrow v \quad E;H \vdash v : T}{E;H \vdash x : T} \quad (\text{T-VARIABLE}) \\ \\ \frac{E;H \vdash a : \{l_1:T_1, \dots, l_n:T_n\} \quad 1 \leq i \leq n}{E;H \vdash a.l_i : T_i} \quad (\text{T-FIELD}) \\ \\ \frac{E;H \vdash e_1 : \text{PrimitiveType} \quad E;H \vdash e_2 : \text{PrimitiveType}}{E;H \vdash e_1 \oplus e_2 : \text{PrimitiveType}} \quad (\text{T-OP}) \\ \\ \frac{E;H \vdash e_1 : T \quad E;H \vdash e_2 : T}{E;H \vdash e_1 \otimes e_2 : \text{PrimitiveType}} \quad (\text{T-COMPARE}) \\ \\ \frac{E;H \vdash e_1 : \text{PrimitiveType} \quad E;H \vdash e_2 : \text{PrimitiveType}}{E;H \vdash e_1 \odot e_2 : \text{PrimitiveType}} \quad (\text{T-COMBINE}) \\ \\ \frac{E;H \vdash e : T}{E;H \vdash \rho e} \quad (\text{T-PRIORITY}) \\ \\ \frac{E;H \vdash C_1 \quad E;H \vdash C_2}{E;H \vdash C_1 \wedge C_2} \quad (\text{T-CONJUNCTION}) \end{array}$$

$$\boxed{E;H \models C}$$

This still represents the call to the solver. The solver now determines values on the heap and in the environment.

$$\boxed{\text{stay}(E) = C}$$
$$\boxed{\text{stay}(H) = C}$$

These judgments define how to translate the environment and heap into a source-level “stay” constraint. The rules are unchanged for primitive values that are stored in the environment, but are amended for the heap. We assume the solver knows about maps to do lookup on the heap.

$$\text{stay}(\emptyset) = \text{true} \quad (\text{STAYEMPTY})$$

$$\frac{\text{E}(\mathbf{x}) = \mathbf{v} \quad \text{E}_0 = \text{E} \setminus \{(\mathbf{x}, \mathbf{v})\} \quad \text{stay}(\text{E}_0) = \text{C}_0 \quad \text{C} = \text{C}_0 \wedge \text{weak } \mathbf{x}=\mathbf{v}}{\text{stay}(\text{E}) = \text{C}} \quad (\text{STAYONE})$$

$$\frac{\mathbf{o} = \{\mathbf{l}_0:\mathbf{v}_0, \dots, \mathbf{l}_n:\mathbf{v}_n\} \quad \text{H}(\mathbf{r}) = \mathbf{o} \quad \text{H}_0 = \text{H} \setminus \{(\mathbf{r}, \mathbf{o})\} \quad \text{stay}(\text{H}_0) = \text{C}_0}{\text{stay}(\text{H}) = \text{C}_0 \wedge \text{H}(\mathbf{r}) = \{\mathbf{l}_0:\mathbf{x}_{r0}, \dots, \mathbf{l}_n:\mathbf{x}_{rn}\} \wedge \text{weak } \mathbf{x}_{r0} = \mathbf{v}_0 \wedge \dots \wedge \text{weak } \mathbf{x}_{rn} = \mathbf{v}_n} \quad (\text{STAYHEAP})$$

Note that the constraints involving $\text{H}(\mathbf{r})$ are *required*, which ensures that a reference keeps pointing at the same object it originally pointed at.

$$\boxed{\langle \text{E} | \text{H} | \text{C} | \text{I} | \mathbf{s} \rangle \longrightarrow \langle \text{E}' | \text{H}' | \text{C}' | \text{I}' \rangle}$$

The semantics for executing statements is extended from what we had before for Babelsberg/Records. When updating a variable, we now solve first for the identity constraints and then for the value constraints. This propagates any type changes through the identity constraints. We then type the value constraints against the environment returned by solving for the identities.

$$\frac{\text{E}; \text{H} \vdash \mathbf{e} \Downarrow \mathbf{v} \quad \text{stay}(\text{E}) = \text{C}_{E_s} \quad \text{stay}(\text{H}) = \text{C}_{H_s} \quad \text{E}'; \text{H}' \models (\text{I} \wedge \text{C}_{E_s} \wedge \text{C}_{H_s} \wedge \mathbf{a}=\mathbf{v}) \quad \text{E}'; \text{H}' \vdash \text{C} \quad \text{E}''; \text{H}'' \models (\text{C} \wedge \text{C}_{E'_s} \wedge \text{C}_{H'_s})}{\langle \text{E} | \text{H} | \text{C} | \text{I} | \mathbf{a} := \mathbf{e} \rangle \longrightarrow \langle \text{E}'' | \text{H}'' | \text{C} | \text{I} \rangle} \quad (\text{S-ASGN})$$

Implicitly this rule gets stuck if either a) the identity constraints cannot be solved, b) the value constraints do not type if the identities are solved, or c) the value constraints cannot be solved. A practical implementation would add explicit exceptions for these cases that the programmer could handle.

$$\frac{\text{E}; \text{H} \vdash \mathbf{e}_1 \Downarrow \mathbf{v}_1 \cdots \text{E}; \text{H} \vdash \mathbf{e}_n \Downarrow \mathbf{v}_n \quad \text{H}(\mathbf{r}) \uparrow \quad \text{H}_0 = \text{H} \cup \{(\mathbf{r}, \{\mathbf{l}_1:\mathbf{v}_1, \dots, \mathbf{l}_n:\mathbf{v}_n\})\} \quad \text{stay}(\text{E}) = \text{C}_{E_s} \quad \text{stay}(\text{H}_0) = \text{C}_{H_s} \quad \text{E}'; \text{H}' \models (\text{I} \wedge \text{C}_{E_s} \wedge \text{C}_{H_s} \wedge \mathbf{x}=\mathbf{r}) \quad \text{E}'; \text{H}' \vdash \text{C} \quad \text{E}''; \text{H}'' \models (\text{C} \wedge \text{C}_{E'_s} \wedge \text{C}_{H'_s})}{\langle \text{E} | \text{H} | \text{C} | \text{I} | \mathbf{x} := \text{new } \{\mathbf{l}_1:\mathbf{e}_1, \dots, \mathbf{l}_n:\mathbf{e}_n\} \rangle \longrightarrow \langle \text{E}'' | \text{H}'' | \text{C} | \text{I} \rangle} \quad (\text{S-ASGNNEW})$$

The next rules are essentially unchanged from their previous versions and just augmented with the heap references.

$$\frac{\text{E}; \text{H} \vdash \text{C}_0 \quad \text{stay}(\text{E}) = \text{C}_{E_s} \quad \text{stay}(\text{H}) = \text{C}_{H_s} \quad \text{E}'; \text{H}' \models (\text{C} \wedge \text{C}_{E_s} \wedge \text{C}_{H_s} \wedge \text{C}_0)}{\langle \text{E} | \text{H} | \text{C} | \text{I} | \text{once } \text{C}_0 \rangle \longrightarrow \langle \text{E}' | \text{H}' | \text{C} | \text{I} \rangle} \quad (\text{S-ONCE})$$

$$\frac{\langle \text{E} | \text{H} | \text{C} | \text{I} | \text{once } \text{C}_0 \rangle \longrightarrow \langle \text{E}' | \text{H}' | \text{C} | \text{I} \rangle \quad \text{C}' = \text{C} \wedge \text{C}_0}{\langle \text{E} | \text{H} | \text{C} | \text{I} | \text{always } \text{C}_0 \rangle \longrightarrow \langle \text{E}' | \text{H}' | \text{C}' | \text{I} \rangle} \quad (\text{S-ALWAYS})$$

$$\langle \text{E} | \text{H} | \text{C} | \text{I} | \text{skip} \rangle \longrightarrow \langle \text{E} | \text{H} | \text{C} | \text{I} \rangle \quad (\text{S-SKIP})$$

$$\frac{\langle \text{E} | \text{H} | \text{C} | \text{I} | \mathbf{s}_1 \rangle \longrightarrow \langle \text{E}' | \text{H}' | \text{C}' | \text{I}' \rangle \quad \langle \text{E}' | \text{H}' | \text{C}' | \text{I}' | \mathbf{s}_2 \rangle \longrightarrow \langle \text{E}'' | \text{H}'' | \text{C}'' | \text{I}'' \rangle}{\langle \text{E} | \text{H} | \text{C} | \text{I} | \mathbf{s}_1; \mathbf{s}_2 \rangle \longrightarrow \langle \text{E}'' | \text{H}'' | \text{C}'' | \text{I}'' \rangle} \quad (\text{S-SEQ})$$

$$\frac{E;H \vdash e \Downarrow \text{true} \quad \langle E|H|C|I|s_1 \rangle \longrightarrow \langle E'|H'|C'|I' \rangle}{\langle E|H|C|I|\text{if } e \text{ then } s_1 \text{ else } s_2 \rangle \longrightarrow \langle E'|H'|C'|I' \rangle} \quad (\text{S-IFTTHEN})$$

$$\frac{E;H \vdash e \Downarrow \text{false} \quad \langle E|H|C|I|s_2 \rangle \longrightarrow \langle E'|H'|C'|I' \rangle}{\langle E|H|C|I|\text{if } e \text{ then } s_1 \text{ else } s_2 \rangle \longrightarrow \langle E'|H'|C'|I' \rangle} \quad (\text{S-IFELSE})$$

$$\frac{E;H \vdash e \Downarrow \text{true} \quad \langle E|C|H|I|s \rangle \longrightarrow \langle E'|H'|C'|I' \rangle \quad \langle E'|H'|C'|I'|\text{while } e \text{ do } s \rangle \longrightarrow \langle E''|H''|C''|I'' \rangle}{\langle E|H|C|I|\text{while } e \text{ do } s \rangle \longrightarrow \langle E''|H''|C''|I'' \rangle} \quad (\text{S-WHILED0})$$

$$\frac{E;H \vdash e \Downarrow \text{false}}{\langle E|H|C|I|\text{while } e \text{ do } s \rangle \longrightarrow \langle E|H|C|I \rangle} \quad (\text{S-WHILESKIP})$$

The next rules are additions to create identity constraints. Note that the above rules for always and once match only constraints that do not include identity comparisons, because those would not type. We require that an identity constraint already be satisfied when it is asserted, as discussed earlier. So the following rules simply need to check the constraint.

$$\frac{E;H \vdash a_0 \Downarrow v \quad E;H \vdash a_1 \Downarrow v}{\langle E|H|C|I|\text{once } a_0 == a_1 \rangle \longrightarrow \langle E|H|C|I \rangle} \quad (\text{S-ONCEIDENTITY})$$

$$\frac{\langle E|H|C|I|\text{once } a_0 == a_1 \rangle \longrightarrow \langle E|H|C|I \rangle \quad I' = I \wedge (a_0 == a_1)}{\langle E|H|C|I|\text{always } a_0 == a_1 \rangle \longrightarrow \langle E|H|C|I' \rangle} \quad (\text{S-ALWAYSIDENTITY})$$

7 Babelsberg/Objects

In this final language, we add support for mutable objects, classes, methods, messages, and inheritance, along with object-oriented constraint definitions (i.e., constraint definitions that can include method calls). This language thus includes all the essential features of actual object constraint languages such as Babelsberg/R, Babelsberg/JS, and Babelsberg/Squeak.

We can thus now write constraints on the results of message sends, such as the following constraint on the x coordinate of the center of a rectangle:

```
always r.center().x() = 100;
```

The constraint here is on the result of sending the message `center` to `r`, then sending the `x` message to that, and finally sending the `=` message to the result. Depending on how the rectangle is stored, the rectangle's center may well be computed in the `center` method rather than simply being looked up; and even `x` might be computed rather than being stored, for example if the point is stored in polar coordinates. Further, expressions such as `a+b` are now treated in an object-oriented fashion, so that this means “send the message `+` with the argument `b` to the object `a`,” with the meaning of `+` in this expression (and in constraints such as `a+b=c`) depending on the class of `a`.

In the earlier Babelsbergs presented in this memo, in the semantics we assume a solver language that is essentially the same as the program language. Constraints that are sent to the constraint solvers could thus be in the program language. This changes for objects. As before, the solver needs to know about uid-records as well as primitive types and i-records; but it does not know about methods or inheritance. So we add a translation in this semantics to inline method calls before passing constraints to the solver.

To accomplish this, we start with a standard model for an object-oriented language with instances, classes, messages, methods, and inheritance, and add constraints on top of that. Some of the classes and methods in the object-oriented language have corresponding primitive types and operations which we can pass directly to the solver. For example, the classes `Integer`, `Float`, and `Boolean` in the source language might be mapped to the primitive integers, floats, and booleans. The semantics includes an automatic boxing and unboxing between instances of these classes and their primitive equivalents. (This boxing and unboxing is part of the formal semantics, but would not necessarily be done in a practical implementation, where we would want primitive types to be represented efficiently.) The `+` method on `Integer` would expand to use the `+` operation on a primitive representation of itself, and similarly for the `+` method for `Float`, the `or` method for `Boolean`, and so forth. However, in Babelsberg/Objects, in contrast to prior languages, we may no longer mix strings and floats, for example. Previously these both typed as `PrimitiveType`, but now these are full objects that can have different structure (e.g., a `floatvalue` field to store the primitive representation on a `Float` object versus a `stringvalue` field on a `String`). While for these cases, the structural compatibility checks could be made to work (e.g., by using a `primvalue` field for these classes), in general there will be classes in the object-oriented language that don't map to a primitive type; and even for a class that does have a corresponding primitive type, that class can have additional methods that aren't mapped to primitive operations. For example, a `factorial` method for `Integer` has no primitive equivalent and has to be translated for the solver.

As with the previous languages, in the semantics, after every statement execution the current set of constraints is solved and the environment is updated. The current set of constraints is determined as follows. If the statement is an assignment, we evaluate the right hand side and then set up a constraint between the resulting value and the left hand side. (As with Babelsberg/UID, and in contrast to the simpler languages, this will be an identity constraint rather than an equality constraint.) Otherwise, it must be a method call, which is evaluated by calling the method normally. In addition to any constraint resulting from an assignment statement, each constraint in the constraint store is added to the set of constraints as well, and the resulting set of constraints is given to the solver. If the solver finds a solution, the environment is updated; but if the constraints have no solution, or if they are too hard for the solver, our execution implicitly gets stuck.

7.1 Value Classes

In the formal semantics, we will use instances of *value classes* as well as instances of ordinary, full-featured, classes. The restrictions on value classes will make it easier to use them with constraints. A number of existing languages have proposals to support forms of value classes, such as Java³ and Scala⁴. In Babelsberg/Objects, instances of value classes are immutable (that is, we need to provide values for all the fields at object creation time, and after that the value class instance cannot be modified).⁵ Finally, object identity is not significant for value class instances — we define the identity test method, but it performs a test for equality rather than identity on its argument. However, value classes are more than simple record declarations, since they support methods and inheritance.

In this section, we will use `Point` and `Rectangle` as our example value classes. To distinguish them from ordinary classes, we will create instances just by supplying the arguments, e.g., `Point(10,20)`, whereas instances of ordinary classes will be created using the `new` message, e.g., `Window.new(...)`.

In a practical implementation of Babelsberg, ideally the host language will itself support value classes, so that we can use them directly. If not, we can use ordinary classes, with appropriate conventions about object

³<http://openjdk.java.net/jeps/169>

⁴<http://docs.scala-lang.org/sips/completed/value-classes.html>

⁵For use with e.g. distributed systems applications, we would also want to restrict instances of value classes to only hold references to primitive types or to other value class, but this restriction isn't needed for our purposes here.

creation, modification, and not testing for object identity.

Finally, to foreshadow the formal semantics, for simplicity we omit details about object creation, method lookup, and inheritance (which are completely standard), and represent ordinary instances as uid-records and value class instances as i-records.

7.2 Control Structures and Methods

Babelsberg/Objects has the simple `if` and `while` control structures that were introduced for Babelsberg/PrimitiveTypes. There are no particular complications in including these in Babelsberg/Objects.

In addition, however, Babelsberg/Objects includes methods. Following our overarching design goal of having a standard object-oriented language in the absence of constraints, in imperative execution mode, methods are standard. Suppose we are evaluating `x.m(a,b)`. We first look in `x`'s class for a method with selector `m`, then its superclass, and so on up the superclass chain. If no such method is found, it is an error.⁶ Otherwise we evaluate the body of the method in a new environment, and return the result. Arguments are passed by value (with pointer semantics), as in most standard object-oriented languages. For instances of primitive types and value classes, the argument can also be copied (since we can't tell the difference between sharing and copying in this case).

Methods called in the ordinary way can create new constraints, which then persist after the method returns. There is an example in the next section to illustrate this.

Methods can also be invoked by constraint expressions. There are a number of restrictions on methods to allow them to be used in this way — briefly, they can't have side effects, and if they consist of more than just a single return statement they can only be used in the forward direction in constraints. Creating an instance of an ordinary class is regarded as a side effect, so methods that can be invoked by constraint expressions can only create instances of value classes. See Section 7.5 for details, including a discussion of the rationale for this particular choice of restrictions.

If the method can be used in a constraint but only in the forward direction, then it can still be called in the usual way, including passing any parameters by value with pointer semantics. For example, suppose we have a constraint `always c=x.m(a,b)`, and method `m` has multiple statements so that it can only be used in the forward direction, i.e., we can find a value for `c` given values for `x`, `a`, and `b` but not any other direction. In that case, `m` can be called in the usual way, passing `a` and `b` by value with pointer semantics. If the other constraints are such that the correct solution involves finding a value for `b` given values for `c`, `x`, and `a`, a practical system will halt with an error that the constraints are too hard for it to solve, due to a method called from the constraint that can't be inlined (or in the formal semantics, we get stuck). However, in contrast to all the previous examples of constraints that were too hard, the fact that they are too hard is a limitation of the transformations we use for methods called by constraints, rather than being a limitation of the solver.

On the other hand, if method `m` can be used multi-directionally, so that for example we can find a value for `b` given values for `x`, `a`, and `c`, then in the semantics the method is inlined so that the resulting constraints can be turned over to the solver, and we can potentially solve for any of `x`, `a`, `b`, or `c`. This done by creating an environment for method `m` and inlining the returned expression with respect to that environment. The parameters in that environment are constrained to be equal to `a` and `b`, and `self` is constrained to be equal to `x`.

⁶This is analogous to the structural compatibility checks from Babelsberg/Records, and just as in Babelsberg/Records we don't invent a new field for a record if needed to satisfy a constraint, in Babelsberg/Objects we don't convert an object to an instance of a different class, or synthesize a new method, to satisfy constraints — instead, if there is a missing method it's just an error. This lookup process isn't represented in our formal semantics since it is standard and doesn't introduce any issues for our focus on constraints.

7.3 Examples

Here is a constraint on the center of a rectangle, where the center is a computed value rather than being stored as an instance variable.

```
always r.center() = Point(10,20);
```

The `center` method for `Rectangle` is defined as follows:

```
def center
  return (self.upper_left + self.lower_right) / 2;
end
```

Note that since `Point` is a value class, we can make an instance of it in the constraint expression itself (“`Point(10,20)`”), and also in the `center` method, since we are doing point addition and division in that method. (If it had been an ordinary class, creating such an instance would have been a side effect, disallowing this constraint.)

The constraint is evaluated in the following way. The code for `Rectangle`’s `center` method is found, and inlined to construct an equality constraint between the result of evaluating `(self.upper_left + self.lower_right) / 2` and a new point with `x=10, y=20`. Constructing the center point from the first expression, as well as looking up and inlining the `=` method on it, explodes into a network of simpler constraints that can then be handed to the solver in one conjunction:

```
// for the center method
self1 = r
self2 = self1.upper_left
arg1 = self1.lower_right
argP1_1 = self2.x + arg1.x
argP1_2 = self2.y + arg1.y
self3 = {x: argP1_1, y: argP2_2}
arg2 = 2
argP2_1 = self3.x / arg2
argP2_2 = self3.y / arg2

// for the static point
argP3_1 = 10
argP3_2 = 20

// for the point equality
self4 = {x: argP2_1, y: argP2_2}
arg3 = {x: argP3_1, y: argP3_2}
self4.x = arg3.x && self4.y = arg3.y
```

For each variable name in a scope a unique global variable name is created. The solver works only on the global names. During evaluation, each local name maps to a global name, which in turn maps to a value. Entering the center method creates a local environment in which `self` is bound to `r`. A global alias `self1` is created and constrained to refer to the same reference (which is to a rectangle). Similarly, the global alias `self2` is a point, namely the upper left corner of the rectangle. This alias is then used when we explode the `+` message to that point. `self3` is also a point, namely the new point that is returned by the point addition

method. Note that we are representing the instance of this value class in the above conjunction of exploded constraints and in the formal semantics as an i-record. For methods that take arguments, global names for the argument names are created and constrained similarly. For example, `arg1` is the global name in this exploded constraint for the argument to the `+` method for `Point`. Thus, the solver never has to deal with different scopes or name clashes.

This example illustrates why we need to solve for object identity before values — if `r` had been an instance of a class with a different implementation of the `center` method, it would have been exploded in a different way.

Here is another example of using a method in a constraint. Suppose we add a `double` method to `Float`, and then use it in a constraint:

```
def double()
  return 2*self;
end;

x := 0;
y := 0;
always y=x.double();
y := 20;
```

After the program runs, `y` will be 20 and `x` will be 10.

For this program, the constraint `y=x.double()` explodes into the following network:

```
self1 = x
y = 2*self1
```

This conjunction of constraints works equally well for determining `x` or determining `y`, even though `double` is actually a message to `x`.

As noted above, methods called in the ordinary way can create new constraints, which then persist after the method returns. Here is a simple example. Suppose we have a `BankAccount` class that includes a `balance` method. We can then define the following method:

```
def require_min_balance(acct,min)
  always acct.balance >= min;
end;
```

Now suppose we call `require_min_balance(a,10)` on some account `a`. (This is just calling the method from an ordinary statement, either at the top level or some other method, *not* from a constraint expression.) Thereafter the balance in the account `a` must be at least 10 (a persistent constraint).

Note the consequences of call by value for such methods. Consider:

```
a := BankAccount.new;
m := 10;
def require_min_balance(a,m);
m := 100;
```

Even though we reassigned `m`, the minimum balance for the account stays at 10, since we passed `m` by value. Of course, we are passing the account `a` by value as well.

```
a := BankAccount.new;
def require_min_balance(a,10);
a := BankAccount.new;
```

After the second assignment, `a` is bound to a new account — but the minimum balance constraint is on the previous instance of `BankAccount` (which will presumably be garbage collected since there aren't any references to it).

If we do want a constraint that continues to be enforced even if the account or the minimum is rebound, we can instead write a method that is a minimum balance test and use it in a constraint:

```
def has_min_balance(acct,min)
  return acct.balance >= min;
end;

a := BankAccount.new;
m := 10
always has_min_balance(a,m);
m := 100;
```

Programmers will need to be aware of the different semantics for these two cases, and select the appropriate variant when it makes a difference.

Implementing such methods requires that the semantics (and implementations) keep the local variables in such a method invocation as long as the constraints it creates are active. Here is a somewhat artificial variant of the original `require_min_balance` method to illustrate this:

```
def require_min_balance_with_locals(acct,min)
  b := acct.balance;
  always b = acct.balance;
  always b >= min;
end;
```

We could also add a minimum balance constraint in the initialization method for `BankAccount`. This version requires that the initial balance be provided, but has a default of 0 for the minimum balance in case it isn't provided.

```
def init(initial_balance,min=0)
  self.balance := initial_balance;
  always self.balance >= min;
end;
```

7.4 Arrays

While it is not strictly necessary to include arrays in Babelsberg/Objects, we have used arrays in some examples that follow. We treat the length of an array is part of its structure, and use the same structural

compatibility checks for the existence of array elements that we do for the existence of fields in a record. Thus, we would not, for example, grow an array to allow a constraint on its i^{th} element to be satisfied if it didn't already have an i^{th} element.

7.5 Additional Restrictions on Constraint Expressions

Methods can of course have side effects, but such methods can't be used in the expressions that define constraints. (This is one of the restrictions noted in Section 4.1 on requirements for constraint expressions.) This subsection describes some consequences of this restriction when we have objects, methods, and object identity. There are also some additional restrictions on methods called from constraints that are used in other than the forward direction, discussed below.

Regarding side effects in methods, consider the `pop` method for a class `Stack`, which has a side effect as well as returning the value popped from the stack. This program fragment is OK:

```
x := stack.pop;
```

This works, because the RHS of the assignment is evaluated first, and then we set up a `once` constraint using the value popped from the stack. (Thus the RHS of an assignment can have side effects, in contrast to `always` or `once` constraints.)

In contrast, these program fragments are not OK:

```
once x == stack.pop;
```

```
always x == stack.pop;
```

It's easy to see why we don't want to allow this: we should be able to evaluate the constraint expression as a test, and if it evaluates to true the constraint is satisfied. This would just not work in these stack examples.

As noted previously, creating an instance of an ordinary class is regarded as a side effect, so methods that can be invoked by constraint expressions can only create instances of value classes and not ordinary classes.⁷

Creating a constraint is a kind of side effect, at least potentially, so another consequence of the prohibition on side effects in constraint expressions is that methods called from a constraint expression cannot themselves create other constraints, or call further methods that do so. The following program is not allowed, and illustrates side effects resulting from adding constraints in methods called from other constraints.

```
def test(i)
  always medium i=5;
  return i+1;
end;
x := 0;
y := 0;
always medium x=10;
always y=test(x);
```

⁷As also noted previously (Footnote 1), in a practical implementation the programmer might be able to make cautious use of benign side effects in a constraint expression. Such benign side effects might include creating instances of ordinary classes, for example, constructing a temporary instance of an ordinary class that is garbage collected before it is visible outside the constraint. In the formal semantics, however, we simply disallow side effects in constraint expressions.

Suppose we are using a least-squares solver that supports soft constraints. Every time we evaluate the expression `y=test(x)`, we add another medium constraint that `x=5`. Since we are finding a least squares solution, this nudges `x` more toward 5 each time, diminishing the influence of the `x=10` constraint.⁸

Here are two other issues with constraints that call methods that add constraints, for the record in case we do want to try to allow this in some restricted set of circumstances:

First, we would need to solve their constraints eagerly rather than just accumulating the constraints and solving all at once. Here's an example that demonstrates this.

```
def test(x)
  always y=x;
  if y>10 then return 1 else return 0;
end;
a := 5;
always b=test(a);
```

If we don't solve eagerly, `y` doesn't have a value when we evaluate the test in the `if` statement.

Also the method could add soft constraints, and itself be invoked by a soft constraint — for example, suppose that in the above program we instead had `always weak b=test(a)`. We would need an algebra for combining priorities, for example, take the minimum of the priorities.

Methods called by constraints can assign to local variables, but not to instance variables or globals (which would be side effects visible outside the method). If there are assignments to local variables, the method can only be used in the forward direction in constraints that call it (i.e., to compute the result given the inputs). The only way that statements in a method other than the final return aren't dead code is if they have side effects, so for practical purposes we can restate this restriction as follows: for a method to be used in other than just the forward direction in constraints that call it, the method must consist of a single `return` statement.

It would be possible to relax this restriction somewhat, so that other kinds of methods could be used in reverse in constraints, but the resulting restrictions would be harder to express, and would have an *ad hoc* feel, likely making it harder for the programmer to keep track of them. So in the design presented here we use a simple, easily understood restriction that covers an important set of practical cases. See Appendix A.5 for a description of a backward compatibility mode for methods that accommodates some methods with multiple statements, which is currently implemented in Babelberg/Ruby.

Here are some examples.

Consider an iterative `sum` method for the class `Array`.

```
def sum
  ans := 0;
  i := 0;
  while i<self.length
    ans := ans + self[i];
    i := i+1;
  end;
  return ans;
end
```

⁸This issue only arises with soft constraints and a global comparator. With constraints that are required, re-adding a required constraint is idempotent. Or if we use a local comparator, re-adding the soft constraint will also not change the resulting solution. But in any case, adding a constraint is conceptually different from testing whether it is satisfied.

This works as expected for normal imperative code, and can also be used in the forward direction in a constraint, for example:

```
a := Array.new(2);
a[0] := 10;
a[1] := 20;
s := 30;
always s = a.sum;
a[0] := 100;
```

After the `always` constraint is executed, `s` is 30; then after the final assignment to `a[0]`, `s` becomes 120.

However, the method doesn't work backwards — for example, we can't constrain the sum of the array and expect the system to update one or more elements to satisfy the constraint. So the constraint in the last line below will be too hard for the system to solve:

```
a := Array.new(2);
a[0] := 10;
a[1] := 20;
always 50 = a.sum;
```

We *can* provide something analogous to an array `sum` constraint that works both forward and backward by writing an ordinary method that sets up the appropriate network of addition constraints. It is defined here as the `sum_equal` method on `Array`, although it could be a standalone method instead.

```
def sum_equal(sum)
  helper(0,sum);
end;

def helper(start,sum)
  if start>=self.length then always sum=0;
  else
    partial := 0; /* create a local variable */
    helper(start+1,partial);
    always sum=self[start]+partial;
  end if;
end;
```

Now our example works:

```
a := Array.new(2);
a[0] := 10;
a[1] := 20;
a.sum_equal(50);
```

If we subsequently set `a[1]` to 5, `a[0]` will become 45 to keep the constraints satisfied.

Note that in contrast to the earlier version, the constraints are not automatically re-applied if `a` is reassigned; the programmer needs to call `sum_equal` again.⁹

It is tempting to write `sum` as a method that returns the sum and that still can be used multi-directionally, perhaps:

```
def sum
  return self.sum_from(0);
end;

def sum_from(start)
  if start >= self.length
    then return 0
    else return self[start] + self.sum_from(start+1);
  end;
end;
```

Unfortunately, this doesn't meet the restriction that methods called from constraints consist only of a single return statement. Worse, if we tried inlining it, the method would expand infinitely. So, at least for now, `sum_equal` seems like it is as good as we can do.

A related example is a method for `Array` that tests whether the argument is one of the elements of that array.

```
def contains(x)
  i:=0;
  while i<self.length && self[i]!=x
    i:=i+1;
  end;
  return i<self.length;
end;
```

Since this includes more than one statement, it can only be used in a constraint in the forward direction:

```
a := Array.new(10);
a.fill(100);
c := false;
always c = a.contains(5);
a[3] := 5;
```

The analogous method to `sum_equal`, say `must_contain`, unfortunately would not be very useful: since we don't have Prolog-style backtracking, it would probably just be satisfied in the reverse direction by setting the first element of the array.

⁹At the cost of complicating the language, we could address this issue by adding a construct that automatically invokes `some_method` whenever needed to re-establish the constraint network. To be used in this way, `some_method` could only assert constraints; it could have no other side effects. There'd also need to be machinery to retract the old constraints before re-running `some_method`, to avoid accumulating obsolete constraints.

7.6 Identity Constraint Examples

The following example illustrates how assignment for objects with object identity is handled. This uses an example class `MutablePoint`, which is an ordinary class whose instances can be changed and that have individual identities, in contrast to the value class `Point` used previously.

```
p1 := MutablePoint.new(10,10);
p2 := p1;
p1 := MutablePoint.new(50,50);
```

The result is just the same as in Ruby or some other object-oriented language without constraints. After the first statement, we have the new point with $x=10, y=10$ in the heap, and solve a required identity constraint that `p1`'s value equal the reference to this point. After the second statement, the identity constraints are a weak identity stay on `p1`, and a required constraint that `p2` reference the same value as `p1`. After the third statement, there is another point with $x=50, y=50$ in the heap, a required constraint that `p1`'s value now equal the reference to this new point, and weak identity constraints that `p1` refer to what it used to (which can't be satisfied), and that `p2` refer to what it used to (which *can* be satisfied.)

As with Babelsberg/UID, in the interests of predictability, the system will not create a new object to satisfy constraints. Thus, this program halts with an error:

```
p := MutablePoint.new(0,0);
q := p;
always p.x=5;
always q.x=10;
```

After the third statement, both `p` and `q` refer to the same point, which has $x=5, y=0$. Trying to satisfy the final constraint on `q.x` fails. We could satisfy the constraints by first cloning `q` from `p` and then changing its `x`, but we forbid creating new objects to satisfy constraints. In the semantics, this results in getting stuck due to an unsatisfiable required constraint.

While cloning a mutable point seems innocuous, consider a similar example with windows on the user's display, which again results in an error. (It seems dangerous to silently create a second window.)

```
w1 := Window.new(...);
w2 := w1;
always w1.width = 200;
always w2.width = 300;
```

However, this variant, in which we make the constraints on the window's widths strong rather than required, is OK. (We don't create a new window to satisfy the strong constraints on the widths.)

```
w1 := Window.new(...);
w2 := w1;
always strong w1.width = 200;
always strong w2.width = 300;
```

As a point of comparison, this restriction on not creating new objects isn't relevant for Babelsberg/Records or when we use instances of value classes in Babelsberg/Objects. Here is one of the above examples using value classes:

```

p := Point(10,20);
q := p;
always p.x=5;
always q.x=10;

```

This is fine, and at the end of the program, `p` and `q` will refer to two different value class instances.

There can be explicit identity constraints, as in Babelberg/UID. Again as in Babelberg/UID, identity constraints must be satisfied when they are first created. For example:

```

x:= Window.new(...);
y := x;
always y==x;
x := Circle.new(...);

```

Here `x` and `y` are identical at the time the identity constraint `always y==x` is created. And given this identity constraint, after the last assignment both `x` and `y` refer to the newly created circle. This example also illustrates how type changes must originate with an assignment statement, but can ripple further via explicit identity constraints.

7.7 Formalism

7.7.1 Syntax

The syntax is augmented to support method invocation as expressions. We introduce syntax for the method body. Additionally, we allow creating new objects as expressions now. We also have a form of immutable records and consider them as “value objects,” which can respond to methods. Finally, we introduce the `nil` value. The entire syntax is given below.

Statement	s	::=	<code>skip</code> <code>L := e</code> <code>always C</code> <code>once C</code> <code>s;s</code> <code>if e then s else s</code> <code>while e do s</code>
Constraint	C	::=	ρe $C \wedge C$
Expression	e	::=	v L $e \oplus e$ $e \otimes e$ $e \odot e$ I $e.l(e_1, \dots, e_n)$ <code>o</code> <code>new o</code> D
Identity	I	::=	$e == e$
Object Literal	o	::=	$\{l_1:e_1, \dots, l_n:e_n\}$
L-Value	L	::=	x $e.l$
Constant	c	::=	<code>true</code> <code>false</code> <code>nil</code> base type constants
Variable	x	::=	variable names
Label	l	::=	record label names
Reference	r	::=	references to heap records
Dereference	D	::=	$\mathbb{H}(r)$
Method Body	b	::=	<code>s; return e</code> <code>return e</code>
Value	v	::=	c r $\{l_1:v_1, \dots, l_n:v_n\}$

We assume that the set of variable names ranged over by metavariable `x` includes the name `self`.

7.7.2 Semantics

Since we have methods with local scopes now, we introduce a global \mathbb{E} that maps global variable names to values. The rules for execution still work on a local environment, but a local environment is only a mapping from local names to globally unique names now. Method lookup creates new local environments, and constraints are translated to translate local variable names to global variable names for the solver. The solver still returns a global environment, and does not know about the local environments and their mapping into the global environment. In addition, the constraint stores \mathbb{I} and \mathbb{C} are replaced with \mathbb{H} and \mathbb{C} , respectively. These now store constraints in pairs with the local environment they were created in. For readability, we now write \mathbb{H} instead of \mathbb{H} , so all the global stores are written in the same font.

$$\boxed{\text{lookup}(\mathbf{v}, \mathbf{l}) = (\mathbf{x}_1 \cdots \mathbf{x}_n, \mathbf{b})}$$

“Lookup of method \mathbf{l} in the object or value \mathbf{v} returns the formal parameter names \mathbf{x}_1 through \mathbf{x}_n and the method body \mathbf{b} ”

This judgment is opaque: our semantics does not depend on how method lookup is performed.

$$\boxed{\text{enter}(\mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, \mathbf{v}, \mathbf{x}_1 \cdots \mathbf{x}_n, \mathbf{e}_1 \cdots \mathbf{e}_n) = (\mathbb{E}', \mathbb{E}_m, \mathbb{H}', \mathbb{C}', \mathbb{I}')}$$

“Invoking a method on \mathbf{v} with argument names \mathbf{x}_1 through \mathbf{x}_n and arguments \mathbf{e}_1 through \mathbf{e}_n constructs the method scope \mathbb{E}_m and may update the heap and constraint stores.”

This is a helper judgment that simplifies the definition of evaluation for method calls (shown below).

$$\frac{\begin{array}{c} \langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_1 | \mathbb{H}_1 | \mathbb{C}_1 | \mathbb{I}_1 | \mathbf{v}_1 \rangle \\ \dots \\ \langle \mathbb{E}_{n-1} | \mathbb{E} | \mathbb{H}_{n-1} | \mathbb{C}_{n-1} | \mathbb{I}_{n-1} | \mathbf{e}_n \rangle \Downarrow \langle \mathbb{E}_n | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n | \mathbf{v}_n \rangle \\ \mathbb{E}_m = \emptyset \quad \langle \mathbb{E}_n | \mathbb{E}_m | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n | \mathbf{self} := \mathbf{v} \rangle \longrightarrow \langle \mathbb{E}_0 | \mathbb{E}_0 | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n \rangle \\ \langle \mathbb{E}_0 | \mathbb{E}_0 | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n | \mathbf{x}_1 := \mathbf{v}_1 \rangle \longrightarrow \langle \mathbb{E}_{n+1} | \mathbb{E}_1 | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n \rangle \\ \dots \\ \langle \mathbb{E}_{2n-1} | \mathbb{E}_{n-1} | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n | \mathbf{x}_n := \mathbf{v}_n \rangle \longrightarrow \langle \mathbb{E}_{2n} | \mathbb{E}_n | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n \rangle \end{array}}{\text{enter}(\mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, \mathbf{v}, \mathbf{x}_1 \cdots \mathbf{x}_n, \mathbf{e}_1 \cdots \mathbf{e}_n) = (\mathbb{E}_{2n}, \mathbb{E}_n, \mathbb{H}_n, \mathbb{C}_n, \mathbb{I}_n)} \quad (\text{ENTER})$$

$$\boxed{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v} \rangle}$$

“Evaluating expression \mathbf{e} produces the value \mathbf{v} , while possibly having side-effects on everything but the local environment.”

$$\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{c} \rangle \Downarrow \langle \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{c} \rangle \quad (\text{E-CONST})$$

$$\frac{\mathbb{E}(\mathbf{x}) = \mathbf{x}_g \quad \mathbb{E}(\mathbf{x}_g) = \mathbf{v}}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{x} \rangle \Downarrow \langle \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{v} \rangle} \quad (\text{E-VAR})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{r} \rangle \quad \mathbb{H}'(\mathbf{r}) = \{\mathbf{l}_1 : \mathbf{v}_1, \dots, \mathbf{l}_n : \mathbf{v}_n\} \quad 1 \leq i \leq n}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}.\mathbf{l}_i \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_i \rangle} \quad (\text{E-FIELD})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \{\mathbf{l}_1 : \mathbf{v}_1, \dots, \mathbf{l}_n : \mathbf{v}_n\} \rangle \quad 1 \leq i \leq n}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}.\mathbf{l}_i \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_i \rangle} \quad (\text{E-VALUEFIELD})$$

$$\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{r} \rangle \Downarrow \langle \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{r} \rangle \quad (\text{E-REF})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_0 | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{v}_1 \rangle \quad \langle \mathbb{E}_0 | \mathbb{E} | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_2 \rangle \quad \mathbf{v}_1 \llbracket \oplus \rrbracket \mathbf{v}_2 = \mathbf{v}}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \oplus \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v} \rangle} \quad (\text{E-OP})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_0 | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{v}_1 \rangle \quad \langle \mathbb{E}_0 | \mathbb{E} | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_2 \rangle \quad \mathbf{v}_1 \llbracket \otimes \rrbracket \mathbf{v}_2 = \mathbf{v}}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \otimes \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v} \rangle} \quad (\text{E-COMPARE})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_0 | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{v}_1 \rangle \quad \langle \mathbb{E}_0 | \mathbb{E} | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_2 \rangle \quad \mathbf{v}_1 \llbracket \odot \rrbracket \mathbf{v}_2 = \mathbf{v}}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \odot \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v} \rangle} \quad (\text{E-COMBINE})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_0 | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{v} \rangle \quad \langle \mathbb{E} | \mathbb{E} | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v} \rangle}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 == \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{true} \rangle} \quad (\text{E-IDENTITYTRUE})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_0 | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{v}_1 \rangle \quad \langle \mathbb{E} | \mathbb{E} | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_2 \rangle \quad \mathbf{v}_1 \neq \mathbf{v}_2}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 == \mathbf{e}_2 \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{false} \rangle} \quad (\text{E-IDENTITYFALSE})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}_0 | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{v} \rangle \quad \text{lookup}(\mathbf{v}, \mathbf{l}) = (\mathbf{x}_1 \cdots \mathbf{x}_n, \mathbf{s}; \text{return } \mathbf{e})}{\langle \mathbb{E}_1 | \mathbb{E}_m | \mathbb{H}_1 | \mathbb{C}_1 | \mathbb{I}_1 | \mathbf{s} \rangle \longrightarrow \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' \rangle \quad \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' | \mathbf{v}_r \rangle} \quad (\text{E-CALL})$$

$$\frac{}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}.1(\mathbf{e}_1, \dots, \mathbf{e}_n) \rangle \Downarrow \langle \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' | \mathbf{v}_r \rangle}$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}_0 | \mathbb{H}_0 | \mathbb{C}_0 | \mathbb{I}_0 | \mathbf{v} \rangle \quad \text{lookup}(\mathbf{v}, \mathbf{l}) = (\mathbf{x}_1 \cdots \mathbf{x}_n, \text{return } \mathbf{e})}{\langle \mathbb{E}_1 | \mathbb{E}_m | \mathbb{H}_1 | \mathbb{C}_1 | \mathbb{I}_1 | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_r \rangle} \quad (\text{E-CALLSIMPLE})$$

$$\frac{}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}.1(\mathbf{e}_1, \dots, \mathbf{e}_n) \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{v}_r \rangle}$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_1 | \mathbb{H}_1 | \mathbb{C}_1 | \mathbb{I}_1 | \mathbf{v}_1 \rangle \cdots \langle \mathbb{E}_{n-1} | \mathbb{E} | \mathbb{H}_{n-1} | \mathbb{C}_{n-1} | \mathbb{I}_{n-1} | \mathbf{e}_n \rangle \Downarrow \langle \mathbb{E}_n | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n | \mathbf{v}_n \rangle}{\mathbb{H}_n(\mathbf{r}) \uparrow \quad \mathbb{H}' = (\mathbb{H}_n \cup \{(\mathbf{r}, \{\mathbf{l}_1 : \mathbf{v}_1, \dots, \mathbf{l}_n : \mathbf{v}_n\})\})} \quad (\text{E-NEW})$$

$$\frac{}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{new} \{ \mathbf{l}_1 : \mathbf{e}_1, \dots, \mathbf{l}_n : \mathbf{e}_n \} \rangle \Downarrow \langle \mathbb{E}_n | \mathbb{H}' | \mathbb{C}_n | \mathbb{I}_n | \mathbf{r} \rangle}$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 \rangle \Downarrow \langle \mathbb{E}_1 | \mathbb{H}_1 | \mathbb{C}_1 | \mathbb{I}_1 | \mathbf{v}_1 \rangle \cdots \langle \mathbb{E}_{n-1} | \mathbb{E} | \mathbb{H}_{n-1} | \mathbb{C}_{n-1} | \mathbb{I}_{n-1} | \mathbf{e}_n \rangle \Downarrow \langle \mathbb{E}_n | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n | \mathbf{v}_n \rangle}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \{ \mathbf{l}_1 : \mathbf{e}_1, \dots, \mathbf{l}_n : \mathbf{e}_n \} \rangle \Downarrow \langle \mathbb{E}_n | \mathbb{H}_n | \mathbb{C}_n | \mathbb{I}_n | \{ \mathbf{l}_1 : \mathbf{v}_1, \dots, \mathbf{l}_n : \mathbf{v}_n \} \rangle} \quad (\text{E-VALUE})$$

$$\boxed{\mathbb{E}; \mathbb{H} \vdash \mathbf{e} : \mathbf{T}}$$

$$\boxed{\mathbb{E}; \mathbb{H} \vdash \mathbf{C}}$$

We typecheck with respect to the global environments \mathbb{E} . Method calls do not typecheck: even though we allow method calls in expressions and thus in constraints syntactically, our rules for creating constraints given below inline method invocations. Identity constraints also do not typecheck: they are solved separately

and should not appear in ordinary constraints (see the rules for statement evaluation below). Finally, new (non-value) object construction does not typecheck, since constraints must be side-effect-free.

Types distinguish between primitive values and objects, and for objects the type keeps track of their fields:

Type $T ::= \text{PrimitiveType} \mid \{l_1:T_1, \dots, l_n:T_n\}$

$$\frac{}{\mathbb{E}; \mathbb{H} \vdash c : \text{PrimitiveType}} \quad (\text{T-CONSTANT})$$

$$\frac{\mathbb{E}(x) = v \quad \mathbb{E}; \mathbb{H} \vdash v : T}{\mathbb{E}; \mathbb{H} \vdash x : T} \quad (\text{T-VARIABLE})$$

We now transform all local variables to their global names using the inlining rules given below before passing them to the solver. Thus, only global variable names type.

$$\frac{\mathbb{E}; \mathbb{H} \vdash e : \{l_1:T_1, \dots, l_n:T_n\} \quad 1 \leq i \leq n}{\mathbb{E}; \mathbb{H} \vdash e.l_i : T_i} \quad (\text{T-FIELD})$$

$$\frac{\mathbb{H}(r)=o \quad \mathbb{E}; \mathbb{H} \vdash o : T}{\mathbb{E}; \mathbb{H} \vdash r : T} \quad (\text{T-REF})$$

$$\frac{\mathbb{E}; \mathbb{H} \vdash e_1 : \text{PrimitiveType} \quad \mathbb{E}; \mathbb{H} \vdash e_2 : \text{PrimitiveType}}{\mathbb{E}; \mathbb{H} \vdash e_1 \oplus e_2 : \text{PrimitiveType}} \quad (\text{T-OP})$$

$$\frac{\mathbb{E}; \mathbb{H} \vdash e_1 : \text{PrimitiveType} \quad \mathbb{E}; \mathbb{H} \vdash e_2 : \text{PrimitiveType}}{\mathbb{E}; \mathbb{H} \vdash e_1 \otimes e_2 : \text{PrimitiveType}} \quad (\text{T-COMPARE})$$

$$\frac{\mathbb{E}; \mathbb{H} \vdash e_1 : \text{PrimitiveType} \quad \mathbb{E}; \mathbb{H} \vdash e_2 : \text{PrimitiveType}}{\mathbb{E}; \mathbb{H} \vdash e_1 \odot e_2 : \text{PrimitiveType}} \quad (\text{T-COMBINE})$$

$$\frac{\mathbb{E}; \mathbb{H} \vdash e_1 : T_1 \cdots \mathbb{E}; \mathbb{H} \vdash e_n : T_n}{\mathbb{E}; \mathbb{H} \vdash \{l_1:e_1, \dots, l_n:e_n\} : \{l_1:T_1, \dots, l_n:T_n\}} \quad (\text{T-VALUEOBJECT})$$

$$\frac{\mathbb{E}; \mathbb{H} \vdash e : T}{\mathbb{E}; \mathbb{H} \vdash \rho e} \quad (\text{T-CONSTRAINT})$$

$$\frac{\mathbb{E}; \mathbb{H} \vdash C_1 : T \quad \mathbb{E}; \mathbb{H} \vdash C_2 : T}{\mathbb{E}; \mathbb{H} \vdash C_1 \wedge C_2} \quad (\text{T-COMPOUNDCONSTRAINT})$$

$$\boxed{\mathbb{E}; \mathbb{H} \models C}$$

As before we assume that the solver natively supports records and uninterpreted functions (which we use to represent the heap). We do not assume, however, that the solver understands methods, which can now be part of constraint expressions. This requires us to essentially inline methods before passing constraints to the solver.

$$\boxed{\text{stay}(\mathbb{E}) = C}$$

$$\boxed{\text{stay}(\mathbb{H}) = \mathbb{C}}$$

The rules for creating stay constraints are adapted to work with multiple environments.

$$\text{stay}(\emptyset) = \text{true} \quad (\text{STAYEMPTY})$$

$$\frac{\mathbb{E}(x) = v \quad \mathbb{E}_0 = \mathbb{E} \setminus \{(x, v)\} \quad \text{stay}(\mathbb{E}_0) = \mathbb{C}_0 \quad \mathbb{C} = \mathbb{C}_0 \wedge \text{weak } x=v}{\text{stay}(\mathbb{E}) = \mathbb{C}} \quad (\text{STAYONE})$$

$$\frac{\mathbb{H}(r) = o \quad \mathbb{H}_0 = \mathbb{H} \setminus \{(r, o)\} \quad \text{stay}(\mathbb{H}_0) = \mathbb{C}_0 \quad o = \{l_0:v_0, \dots, l_n:v_n\}}{\text{stay}(\mathbb{H}) = \mathbb{C}_0 \wedge \mathbb{H}(r) = \{l_0:x_{r0}, \dots, l_n:x_{rn}\} \wedge \text{weak } x_{r0} = v_0 \wedge \dots \wedge \text{weak } x_{rn} = v_n} \quad (\text{STAYHEAP})$$

$$\boxed{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}', e_C, e' \rangle}$$

“Inlining expression e from the local environment \mathbb{E} turns into expression e' . To connect variables across method calls, the constraint expression e_C is returned.”

We use an inlining judgment to translate expressions into a representation suitable for the solver. In particular, we translate local variables into their names in the global environment and provide a semantics for method calls inside constraints. Arguments to method calls are constrained to be equal to the expression that generated them. Inlining does not allow updates to the heap, so no new heap is returned. We do allow assignments to locals in inlined methods, however, so the global environment can change as a result of inlining.

$$\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, c \rangle \rightsquigarrow \langle \mathbb{E}, \text{true}, c \rangle \quad (\text{I-CONST})$$

$$\frac{\mathbb{E}(x) = x_g}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, x \rangle \rightsquigarrow \langle \mathbb{E}, \text{true}, x_g \rangle} \quad (\text{I-VAR})$$

$$\frac{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \rangle \rightsquigarrow \langle \mathbb{E}_1, e_{C_1}, e'_1 \rangle \cdots \langle \mathbb{E}, \mathbb{E}, \mathbb{H}_{n-1}, \mathbb{C}, \mathbb{I}, e_n \rangle \rightsquigarrow \langle \mathbb{E}_n, e_{C_n}, e'_n \rangle}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, \{l_1:e_1, \dots, l_n:e_n\} \rangle \rightsquigarrow \langle \mathbb{E}_n, e_{C_1} \wedge \cdots \wedge e_{C_n}, \{l_1:e'_1, \dots, l_n:e'_n\} \rangle} \quad (\text{I-VALUE})$$

$$\frac{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}', e_C, e' \rangle \quad \langle \mathbb{E}' | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}'' | \mathbb{H} | \mathbb{C} | \mathbb{I} | r \rangle}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.l \rangle \rightsquigarrow \langle \mathbb{E}', e_C, \mathbb{H}(e').l \rangle} \quad (\text{I-FIELD})$$

$$\frac{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}', e_C, e' \rangle \quad \langle \mathbb{E}' | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}'' | \mathbb{H} | \mathbb{C} | \mathbb{I} | \{l_1:v_1, \dots, l_n:v_n\} \rangle}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.l \rangle \rightsquigarrow \langle \mathbb{E}', e_C, e'.l \rangle} \quad (\text{I-VALUEFIELD})$$

$$\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, r \rangle \rightsquigarrow \langle \mathbb{E}, \text{true}, r \rangle \quad (\text{I-REF})$$

$$\frac{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \rangle \rightsquigarrow \langle \mathbb{E}', e_{C_a}, e_a \rangle \quad \langle \mathbb{E}', \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_b}, e_b \rangle}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \oplus e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_a} \wedge e_{C_b}, e_a \oplus e_b \rangle} \quad (\text{I-OP})$$

$$\frac{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \rangle \rightsquigarrow \langle \mathbb{E}', e_{C_a}, e_a \rangle \quad \langle \mathbb{E}', \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_b}, e_b \rangle}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \otimes e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_a} \wedge e_{C_b}, e_a \otimes e_b \rangle} \quad (\text{I-COMPARE})$$

$$\frac{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \rangle \rightsquigarrow \langle \mathbb{E}', e_{C_a}, e_a \rangle \quad \langle \mathbb{E}', \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_b}, e_b \rangle}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \circledast e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_a} \wedge e_{C_b}, e_a \circledast e_b \rangle} \quad (\text{I-COMBINE})$$

$$\frac{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \rangle \rightsquigarrow \langle \mathbb{E}', e_{C_a}, e_a \rangle \quad \langle \mathbb{E}', \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_b}, e_b \rangle}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 == e_2 \rangle \rightsquigarrow \langle \mathbb{E}'', e_{C_a} \wedge e_{C_b}, e_a == e_b \rangle} \quad (\text{I-IDENTITY})$$

$$\frac{\begin{array}{l} \langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H} | \mathbb{C} | \mathbb{I} | v \rangle \quad \text{lookup}(v, l) = (x_1 \cdots x_n, s; \text{return } e) \\ \text{enter}(\mathbb{E}', \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, v, x_1 \cdots x_n, e_1 \cdots e_n) = (\mathbb{E}'', \mathbb{E}_m, \mathbb{H}, \mathbb{C}, \mathbb{I}) \\ \langle \mathbb{E}'' | \mathbb{E}_m | \mathbb{H} | \mathbb{C} | \mathbb{I} | s \rangle \longrightarrow \langle \mathbb{E}''' | \mathbb{E}' | \mathbb{H} | \mathbb{C} | \mathbb{I} \rangle \quad \langle \mathbb{E}''' | \mathbb{E}' | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}'''' | \mathbb{H} | \mathbb{C} | \mathbb{I} | v_r \rangle \end{array}}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.l(e_1, \dots, e_n) \rangle \rightsquigarrow \langle \mathbb{E}'''' , \text{true}, v_r \rangle} \quad (\text{I-CALL})$$

Methods that have any statements at all can only be used in a one-way manner. This is ensured by evaluating the return expression and using only the value in the constraint. Because we are retranslating all constraints on each semantic step, this return value will get updated when its dependencies change, it just won't work in the other direction.

When inlining a method with more than one statement, the statements are simply executed. In particular, this means that we eagerly choose which branch of if-statements to inline and eagerly unroll loops. Further, the I-CALL rule above and the I-MULTIWAYCALL rule below ensure that methods being used in constraints have no side effects. This is accomplished by requiring the initial heap to remain unchanged.

Similarly, methods used in constraints cannot declare nested constraints; this is accomplished by requiring the initial sets of ordinary and identity constraints to remain unchanged.

$$\frac{\begin{array}{l} \langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_0 \rangle \rightsquigarrow \langle \mathbb{E}', e_{C_0}, e'_0 \rangle \quad \langle \mathbb{E}' | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e_0 \rangle \Downarrow \langle \mathbb{E}'' | \mathbb{H} | \mathbb{C} | \mathbb{I} | v \rangle \\ \text{lookup}(v, l) = (x_1 \cdots x_n, \text{return } e) \\ \text{enter}(\mathbb{E}'', \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, v, x_1 \cdots x_n, e_1 \cdots e_n) = (\mathbb{E}''', \mathbb{E}_m, \mathbb{H}, \mathbb{C}, \mathbb{I}) \\ \langle \mathbb{E}''' | \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_1 \rangle \rightsquigarrow \langle \mathbb{E}_1, e_{C_1}, e'_1 \rangle \cdots \langle \mathbb{E}_{n-1}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_n \rangle \rightsquigarrow \langle \mathbb{E}_n, e_{C_n}, e'_n \rangle \\ \mathbb{E}_m(\text{self}) = x_{g_{\text{self}}} \quad \mathbb{E}_m(x_1) = x_{g_1} \cdots \mathbb{E}_m(x_n) = x_{g_n} \\ e_C = (x_{g_{\text{self}}} = e'_0 \wedge x_{g_1} = e'_1 \wedge \cdots \wedge x_{g_n} = e'_n) \quad \langle \mathbb{E}_n, \mathbb{E}_m, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}'_n, e_{C_m}, e' \rangle \end{array}}{\langle \mathbb{E}, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_0.l(e_1, \dots, e_n) \rangle \rightsquigarrow \langle \mathbb{E}'_n, e_C \wedge e_{C_m} \wedge e_{C_0} \wedge e_{C_1} \wedge \cdots \wedge e_{C_n}, e' \rangle} \quad (\text{I-MULTIWAYCALL})$$

For methods that only return an expression, we inline the expression and pass it to the solver. Note that we execute the argument expressions and receiver for their value (potentially executing through other methods), and also inline them, potentially executing the same methods twice (once through I-CALL and once through E-CALL.) Although not ideal in terms of providing the cleanest possible semantics, in practical terms this should not be a problem, because we prohibit side-effects in these methods.

$$\boxed{\langle \mathbb{E}, \mathbb{H}, \mathbb{I}, \mathbb{C} \rangle \rightsquigarrow \langle \mathbb{E}', \mathbb{C} \rangle}$$

$$\boxed{\langle \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I} \rangle \rightsquigarrow \langle \mathbb{E}', \mathbb{C} \rangle}$$

“Re-inlining the constraint store \mathbb{C} returns a constraint \mathbb{C} ”

“Re-inlining the constraint store \mathbb{I} returns a constraint \mathbb{C} ”

$$\langle \mathbb{E}, \mathbb{H}, \mathbb{I}, \emptyset \rangle \rightsquigarrow \langle \mathbb{E}, \text{true} \rangle \quad (\text{I-REINLINEEMPTYC})$$

$$\langle \mathbb{E}, \mathbb{H}, \mathbb{C}, \emptyset \rangle \rightsquigarrow \langle \mathbb{E}, \text{true} \rangle \quad (\text{I-REINLINESKIP})$$

$$\frac{\mathbb{C}_0 = \mathbb{C} \setminus \{(\mathbb{E}, \rho \mathbf{e})\} \quad \langle \mathbb{E}, \mathbb{H}, \mathbb{I}, \mathbb{C}_0 \rangle \rightsquigarrow \langle \mathbb{E}_0, \mathbb{C}_0 \rangle \quad \langle \mathbb{E}_0, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, \mathbf{e} \rangle \rightsquigarrow \langle \mathbb{E}', \mathbf{e}_{C_e}, \mathbf{e}' \rangle}{\langle \mathbb{E}, \mathbb{H}, \mathbb{I}, \mathbb{C} \rangle \rightsquigarrow \langle \mathbb{E}', \mathbb{C}_0 \wedge \rho (\mathbf{e}' \wedge \mathbf{e}_{C_e}) \rangle} \quad (\text{I-REINLINEC})$$

$$\frac{\mathbb{I}_0 = \mathbb{I} \setminus \{(\mathbb{E}, \text{required } \mathbf{e})\} \quad \langle \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}_0 \rangle \rightsquigarrow \langle \mathbb{E}_0, \mathbb{C}_0 \rangle \quad \langle \mathbb{E}_0, \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I}, \mathbf{e} \rangle \rightsquigarrow \langle \mathbb{E}', \mathbf{e}_{C_e}, \mathbf{e}' \rangle}{\langle \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I} \rangle \rightsquigarrow \langle \mathbb{E}', \mathbb{C}_0 \wedge \text{required } (\mathbf{e}' \wedge \mathbf{e}_{C_e}) \rangle} \quad (\text{I-REINLINEI})$$

$$\boxed{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbb{C} \rangle \Longrightarrow \langle \mathbb{E}' | \mathbb{H}' \rangle}$$

This is a helper judgment for use in evaluating assignment statements and identity constraints. These statements are the only ones that can cause the types of variables to change. We handle this in two phases: in a first phase, we propagate the new equality constraint through all existing identity constraints in order to update other variables and fields as needed. In the second phase we solve all of the non-identity constraints, plus any constraints \mathbb{C} created by inlining identity constraints.

$$\frac{\begin{array}{l} \text{stay}(\mathbb{E}) = \mathbb{C}_{E_s} \quad \text{stay}(\mathbb{H}) = \mathbb{C}_{H_s} \quad \langle \mathbb{E}, \mathbb{H}, \mathbb{C}, \mathbb{I} \rangle \rightsquigarrow \langle \mathbb{E}_i, \mathbb{C}_i \rangle \\ \mathbb{E}'; \mathbb{H}' \models (\mathbb{C}_i \wedge \mathbb{C}_{E_s} \wedge \mathbb{C}_{H_s} \wedge \mathbf{e}_1 = \mathbf{e}_2) \\ \text{stay}(\mathbb{E}') = \mathbb{C}_{E'_s} \quad \text{stay}(\mathbb{H}') = \mathbb{C}_{H'_s} \quad \langle \mathbb{E}', \mathbb{H}', \mathbb{I}, \mathbb{C} \rangle \rightsquigarrow \langle \mathbb{E}_c, \mathbb{C} \rangle \quad \mathbb{E}'; \mathbb{H}' \vdash \mathbb{C} \\ \mathbb{E}''; \mathbb{H}'' \models (\mathbb{C} \wedge \mathbb{C}_0 \wedge \mathbb{C}_{E'_s} \wedge \mathbb{C}_{H'_s}) \end{array}}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e}_1 == \mathbf{e}_2 | \mathbb{C}_0 \rangle \Longrightarrow \langle \mathbb{E}'' | \mathbb{H}'' \rangle} \quad (\text{TWO PHASE UPDATE})$$

$$\boxed{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{s} \rangle \longrightarrow \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' \rangle}$$

The stepping rules are refactored to work with the local environments and the new constraint and identity-constraint stores.

$$\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \text{skip} \rangle \longrightarrow \langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} \rangle \quad (\text{S-SKIP})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{s}_1 \rangle \longrightarrow \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' \rangle \quad \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{s}_2 \rangle \longrightarrow \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' \rangle}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{s}_1; \mathbf{s}_2 \rangle \longrightarrow \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' \rangle} \quad (\text{S-SEQ})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \text{true} \rangle \quad \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{s}_1 \rangle \longrightarrow \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' \rangle}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \text{if } \mathbf{e} \text{ then } \mathbf{s}_1 \text{ else } \mathbf{s}_2 \rangle \longrightarrow \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' \rangle} \quad (\text{S-IFTHEN})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \text{false} \rangle \quad \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{s}_2 \rangle \longrightarrow \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' \rangle}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \text{if } \mathbf{e} \text{ then } \mathbf{s}_1 \text{ else } \mathbf{s}_2 \rangle \longrightarrow \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' \rangle} \quad (\text{S-IFELSE})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \text{true} \rangle \quad \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \mathbf{s} \rangle \longrightarrow \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' \rangle \quad \langle \mathbb{E}'' | \mathbb{E}'' | \mathbb{H}'' | \mathbb{C}'' | \mathbb{I}'' | \text{while } \mathbf{e} \text{ do } \mathbf{s} \rangle \longrightarrow \langle \mathbb{E}''' | \mathbb{E}''' | \mathbb{H}''' | \mathbb{C}''' | \mathbb{I}''' \rangle}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \text{while } \mathbf{e} \text{ do } \mathbf{s} \rangle \longrightarrow \langle \mathbb{E}''' | \mathbb{E}''' | \mathbb{H}''' | \mathbb{C}''' | \mathbb{I}''' \rangle} \quad (\text{S-WHILED O})$$

$$\frac{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \mathbf{e} \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' | \text{false} \rangle}{\langle \mathbb{E} | \mathbb{E} | \mathbb{H} | \mathbb{C} | \mathbb{I} | \text{while } \mathbf{e} \text{ do } \mathbf{s} \rangle \longrightarrow \langle \mathbb{E}' | \mathbb{E}' | \mathbb{H}' | \mathbb{C}' | \mathbb{I}' \rangle} \quad (\text{S-WHILESKIP})$$

$$\frac{\begin{array}{c} E(x)\uparrow \quad E(x_g)\uparrow \quad E' = E \cup \{(x, x_g)\} \\ \langle E | E | H | C | I | e \rangle \Downarrow \langle E' | H' | C' | I' | v \rangle \\ \langle E | E' | H | C | I | x_g == v | \text{true} \rangle \Longrightarrow \langle E' | H' \rangle \end{array}}{\langle E | E | H | C | I | x := e \rangle \longrightarrow \langle E' | E' | H' | C | I \rangle} \quad (\text{S-ASGNNEWLOCAL})$$

$$\frac{\begin{array}{c} E(x) = x_g \\ \langle E | E | H | C | I | e \rangle \Downarrow \langle E' | H' | C' | I' | v \rangle \\ \langle E | E' | H | C | I | x_g == v | \text{true} \rangle \Longrightarrow \langle E' | H' \rangle \end{array}}{\langle E | E | H | C | I | x := e \rangle \longrightarrow \langle E' | E' | H' | C | I \rangle} \quad (\text{S-ASGNLOCAL})$$

$$\frac{\begin{array}{c} \langle E | E | H | C | I | e \rangle \Downarrow \langle E' | H' | C' | I' | v \rangle \quad \langle E', E, H', C, I, e_l.1 \rangle \rightsquigarrow \langle E'', \text{true}, e' \rangle \\ \langle E | E | H | C | I | e' == v | \text{true} \rangle \Longrightarrow \langle E' | H' \rangle \end{array}}{\langle E | E | H | C | I | e_l.1 := e \rangle \longrightarrow \langle E' | E | H' | C | I \rangle} \quad (\text{S-ASGNLVALUE})$$

Note that the above rule can only be used in constraint-construction mode if $H=H'$. This effectively restricts assignments in constraints to work on values. Note also, that assignments in constraints are just executed, but do not set up a required equality constraint between the LHS and RHS. (The current Babelberg/Ruby includes a “backwards compatibility mode” which sets up required equality constraints when assignment occurs, but it’s not clear whether we should retain support for that. See Appendix A.5.)

We also do not allow the use of the following rules for **once** and **always** in constraint-construction mode, because the inlining rules disallow updating the constraint store and heap. As discussed previously (Section 7.5), in a practical implementation we might want to support benign side effects in methods that are invoked by constraint expressions, including methods that themselves create new constraints; but this is not modeled by this semantics.

$$\frac{\begin{array}{c} C_0 = \rho e \quad \langle E, E, H, C, I, e \rangle \rightsquigarrow \langle E', e_{C_e}, e' \rangle \quad C'_0 = \rho (e' \wedge e_{C_e}) \quad E'; H \vdash C'_0 \\ \text{stay}(E') = C_{E_s} \quad \text{stay}(H) = C_{H_s} \quad \langle E', H, I, C \rangle \rightsquigarrow \langle E'', C \rangle \quad E'''; H' \models (C \wedge C_{E_s} \wedge C_{H_s} \wedge C'_0) \end{array}}{\langle E | E | H | C | I | \text{once } C_0 \rangle \longrightarrow \langle E''' | E | H' | C | I \rangle} \quad (\text{S-ONCE})$$

$$\frac{\begin{array}{c} \langle E | E | H | C | I | \text{once } C_0 \rangle \longrightarrow \langle E' | E | H' | C | I \rangle \\ C' = C \cup \{(E, C_0)\} \end{array}}{\langle E | E | H | C | I | \text{always } C_0 \rangle \longrightarrow \langle E' | E | H' | C' | I \rangle} \quad (\text{S-ALWAYS})$$

$$\frac{\begin{array}{c} \langle E | E | H | C | I | e_0 \rangle \Downarrow \langle E_0 | H | C | I | v \rangle \quad \langle E_0 | E | H | C | I | e_1 \rangle \Downarrow \langle E_1 | H | C | I | v \rangle \\ \langle E_1, E, H, C, I, e_0 \rangle \rightsquigarrow \langle E_2, e_{C_0}, e'_0 \rangle \quad \langle E_2, E, H, C, I, e_1 \rangle \rightsquigarrow \langle E_3, e_{C_1}, e'_1 \rangle \end{array}}{\langle E | E | H | C | I | \text{once } e_0 == e_1 \rangle \longrightarrow \langle E' | E | H' | C | I \rangle} \quad (\text{S-ONCEIDENTITY})$$

$$\frac{\begin{array}{c} \langle E | E | H | C | I | \text{once } e_0 == e_1 \rangle \longrightarrow \langle E' | E | H' | C | I \rangle \\ I' = I \cup \{(E, e_0 == e_1)\} \end{array}}{\langle E | E | H | C | I | \text{always } e_0 == e_1 \rangle \longrightarrow \langle E' | E | H' | C | I' \rangle} \quad (\text{S-ALWAYSIDENTITY})$$

References

- [1] Edward A. Ashcroft and William W. Wadge. Lucid, a nonprocedural language with iteration. *Communications of the ACM*, 20(7):519–526, July 1977.
- [2] Greg J Badros, Alan Borning, and Peter J Stuckey. The Cassowary linear arithmetic constraint solving algorithm. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 8(4):267–306, 2001.
- [3] Alan Borning, Bjorn Freeman-Benson, and Molly Wilson. Constraint hierarchies. *Lisp and Symbolic Computation*, 5(3):223–270, September 1992.
- [4] Tim Felgentreff, Alan Borning, and Robert Hirschfeld. Babelsberg: Specifying and solving constraints on object behavior. Technical Report 81, Hasso Plattner Institute, University of Potsdam, Potsdam, Germany, May 2014. Also published as TR-2013-001, Viewpoints Research Institute, Los Angeles, CA.
- [5] Tim Felgentreff, Alan Borning, Robert Hirschfeld, Jens Lincke, Yoshiki Ohshima, Bert Freudenberg, and Robert Krahn. Babelsberg/JS: A browser-based implementation of an object constraint language. In *Proceedings of the 2014 European Conference on Object-Oriented Programming*. Springer, July 2014. In press.
- [6] Bjorn Freeman-Benson. Kaleidoscope: Mixing objects, constraints, and imperative programming. In *Proceedings of the 1990 Conference on Object-Oriented Programming Systems, Languages, and Applications, and European Conference on Object-Oriented Programming*, pages 77–88, Ottawa, Canada, October 1990. ACM.
- [7] Bjorn Freeman-Benson. *Constraint Imperative Programming*. PhD thesis, University of Washington, Department of Computer Science and Engineering, July 1991. Published as Department of Computer Science and Engineering Technical Report 91-07-02.
- [8] Bjorn Freeman-Benson and Alan Borning. The design and implementation of Kaleidoscope’90, a constraint imperative programming language. In *Proceedings of the IEEE Computer Society International Conference on Computer Languages*, pages 174–180, April 1992.
- [9] Bjorn Freeman-Benson and Alan Borning. Integrating constraints with an object-oriented language. In *Proceedings of the 1992 European Conference on Object-Oriented Programming*, pages 268–286, June 1992.
- [10] Bjorn Freeman-Benson, John Maloney, and Alan Borning. An incremental constraint solver. *Communications of the ACM*, 33(1):54–63, January 1990.
- [11] Joxan Jaffar and Jean-Louis Lassez. Constraint logic programming. In *Proceedings of the Fourteenth ACM Principles of Programming Languages Conference*, Munich, January 1987.
- [12] Gus Lopez, Bjorn Freeman-Benson, and Alan Borning. Constraints and object identity. In *Proceedings of the 1994 European Conference on Object-Oriented Programming*, pages 260–279, July 1994.
- [13] Michael J. Maher. Logic Semantics for a Class of Committed-choice Programs. In *Proceedings of the Fourth International Conference on Logic Programming*, pages 858–876, Melbourne, May 1987.

A Appendix

This appendix contains some observations that are potentially of interest but that are mostly orthogonal to the task at hand of providing a formal semantics for Babelsberg; and also descriptions of some alternatives that were considered and later dropped, but that seem promising enough that we might come back to them some day. (In other words, it’s a holding pen for material we probably should just have deleted but couldn’t quite bear to.)

A.1 Warnings and Debugging

We don’t deal with warnings and debugging in the formal semantics for Babelsberg. In a practical, usable implementation, however, including such support will be key.

First, all of the places where the formal semantics gets stuck to model an error state should have explicit, understandable error messages in a real implementation. A few additional issues:

A non-required `always` constraint might undo the result of an assignment.

```
x := 0;
always strong x=10;
x := 5;
```

`x` starts out as 0, then is 10 to satisfy the `always` constraint, then gets set to 5, and then at the next time step, the `always` constraint causes it to be set back to 10. We may want to issue a warning in this case in a practical implementation. (This would be a simple case to check for.)

As described previously, in the interests of predictability, the system will not create a new (standard) object to satisfy constraints. There should be an appropriate error message in such a case if the unsatisfied constraints are required, and a warning if they are soft.

A.2 The Perturbation Model vs. the Refinement Model

A key issue in integrating constraints with imperative programming is how to represent variables that change over time. Earlier versions of Kaleidoscope [7, 8, 9] used a *refinement model*, in which ordinary variables were represented as a stream of “pellucid variables,” each holding a value at a different time. This was in turn adapted from Lucid [1]. Later versions of Kaleidoscope [12] used a *perturbation model*, in which variables were represented conventionally. Assignment perturbed the value of a variable, and then the constraints took over and adjusted the values of other variables so that they were re-satisfied if necessary. The formal semantics described here uses a variation of the perturbation model, which we believe is cleaner. (In particular, rather than just perturbing the value of a variable by changing it, we model assignment as a `once` constraint between the variable and the value of the expression on the right hand side of an assignment statement.)

The refinement model does provide additional capabilities — for example syntax for referring to both the current and previous values of a variable — but at a cost, both conceptually for the programmer and also for the language implementor. At least for the common cases, we believe the refinement model and our current model provide the same answers.

The refinement model makes essential use of read-only annotations on variables to prevent the present from affecting the past. So before launching into an example, we digress to discuss read-only annotations.

Intuitively, when choosing a best solution to a set of constraints, constraints should not be allowed to affect the choice of values for their read-only variables, i.e., information can flow out of the read-only variables, but not into them. There is a formal, declarative definition of read-only annotations in [3], which was in turn adapted from that in the ALPS flat committed-choice logic language [13]. A *one-way constraint* can be represented by annotating all but one of the constrained variables as read-only. One can also annotate an expression as read-only — this is syntactic sugar for annotating each of its variables as read-only.

Here’s a simple example of the effect of a read-only annotation on a variable. First consider these constraints:

```
required  x = y
strong   x = 3
weak     y = 4
```

This has the solution $x \mapsto 3, y \mapsto 3$, since the strong constraint trumps the weak one.

But suppose we add a read-only annotation:

```
required  x = y?
strong   x = 3
weak     y = 4
```

Because y is read-only in the $x = y?$ constraint, the solver can’t use that constraint in determining the value of y , and so the solution this time is $x \mapsto 4, y \mapsto 4$.

Further, if we make the constraint $x = 3$ be required:

```
required  x = y?
required  x = 3
weak     y = 4
```

the constraints now have no solution: on the one hand, any solution must satisfy the required $x = 3$ constraint, but on the other, we still can’t use $x = y?$ in determining a value for y . (This is a so-called “blocked” set of constraints [3].)

With those preliminaries out of the way, consider again the example in Section 4 of unsatisfiable constraints arising from the following program:

```
x := 5;
always x <= 10;
x := x + 15
```

This behavior we modeled with the rules for Babelsberg/PrimitiveTypes is the same as that of the refinement model, in which the program would be equivalent to these constraints:

$$\begin{aligned} & \text{required } x_0 = 5 \\ \forall t > 0 & \text{ required } x_t \leq 10 \\ & \text{weak } x_1 = x_0? \\ & \text{weak } x_2 = x_1? \\ & \text{required } x_2 = x_1? + 15 \end{aligned}$$

The read-only annotations in the refinement model serve the same role as do the evaluation rules in our current model. In the example, in our current semantics we model the final assignment $x := x + 15$ by first

evaluating $x+15$ in the old environment, and then adding a required `once` constraint that x be equal to that value. In the refinement model, we model the final assignment as $x_2 = x_1 + 15$, where the read-only annotation on x_1 accomplishes the same thing, by preventing the solver from changing x_1 to satisfy this constraint even though the other constraint that gives x_1 the value of 5 is only weak. Similarly, stay constraints in the current semantics are modeled as weak constraints equating the variable with its current value; in the refinement model these are weak constraints relating the variables representing the current and previous versions, with the previous version annotated as read-only.

A.3 Issues with Using Once Constraints in the Semantics for Assignments

In all of the Babelsbergs, in the semantics we model an assignment by evaluating the right-hand-side, creating a new variable if needed for the left-hand-side, setting it to the new value, and then solving the set of constraints that include a `once` constraint that the left-hand-side equal that value. This lets assignment interact correctly with other constraints in the constraint store, if any. For example, as a consequence of this semantics, the following program will fail (i.e., get stuck in the formal semantics; or raise an exception in a practical implementation):

```
x := 0;
always x=10;
x := 5;
```

However, there are some subtle issues in connection with this representation.

The analogous program still fails, as one would want, in Babelsberg/Records or when using value classes, since records and instances of value classes are immutable.

```
x := {a:0};
always x.a=10;
x := {a:5};
```

But once we get to Babelsberg/UID and Babelsberg/Objects with ordinary objects, the situation changes. Consider the analogous program in Babelsberg/UID:

```
x := new {a:0};
always x.a=10;
x := new {a:5};
```

Here, the right-hand-side of the last assignment is evaluated, the reference to the new record is assigned to x , and then the system immediately changes its `a` field to 10 — it is only the reference that is immutable, not the record on the heap. This is arguably weird, and so one of the design alternatives we considered was to make the right hand side be read-only (recursively). However, this might be computationally expensive, and would also cause difficulties when creating circular structures (see below). So in the current design, we just evaluate the right-hand-side to a value, but don't do anything further about making it recursively read-only.

A.3.1 Circular Structures

There is an issue regarding creating circular structures if we convert assignment statements to `once` constraints, with the value on the right hand side recursively annotated as read only. (The same issue arose in

an earlier version of the semantics in which we simply turned an assignment into a **once** constraint with the right hand side annotated as read only, instead of first evaluating the right hand side.) Consider:

```
c := Cons.new(10,nil);
c.cdr := c;
```

If the second assignment is converted to the constraint **once** `c.cdr == c?` this is unsatisfiable. However, there is a simple workaround, namely to replace such an assignment with a **once** identity constraint without the read-only annotation:

```
c := Cons.new(10,nil);
once c.cdr == c
```

This would be a bit strange. But this solution doesn't require any changes to the formal semantics, and addresses all the previous issues.

Selectively Eliding the Read-Only Annotation. A variant that provides a more standard syntax is to make the entire object on the right-hand side of the assignment be read-only, except for a field that is assigned to (if any). The constraints would be exactly the same for all the examples except for the circular structures one:

```
c := Cons.new(10,nil);
c.cdr := c;
```

Here, `c.car` and the reference to `c` itself are read-only, but not `c.cdr`.

Distinguishing Ownership from Reference. Another alternative to making the entire object on the right-hand side of the assignment read-only is to distinguish ownership of parts from references to other objects. If an instance variable is for an “owned” part, then that instance variable is made read-only; but otherwise not. For compatibility with the host language, the default would be that instance variables do not refer to owned parts; rather, this must be declared explicitly. (It only needs to be declared if there are constraints on the parts or subparts.)

A.4 Alternatives to Structural Compatibility Checks

Our model includes structural compatibility checks on constraints, which prevent the solver from potentially generating new kinds of records and other odd behaviors. But since these ideas seem to keep resurfacing, we outline some of our rejected alternatives in this subsection, in case they end up being useful later.

We first consider records in Babelsberg/Records (immutable records, no UUIDs).

If the solver will need to compare two records with different fields, we need to extend the comparators to handle them. For LPB, there are at least two possibilities:

1. An equality constraint between records is either satisfied or it's not. For example consider a constraint $p=q$ when $p=\{x:0, y:1\}$ and $q=\{x:0, y:2\}$, versus when $q=\{a:1000\}$. In both cases the constraint is unsatisfied, and there is no reason to prefer one of these solutions over the other.
2. An equality constraint between records is unfolded into multiple constraints on the fields, each of which is satisfied or not. (Such a constraint might be unsatisfied either because the values in one or

more corresponding fields weren't equal, or were of different types, or because the corresponding field didn't even exist.) In the above example, for the first case we have $p.x=q.x$ (satisfied) and $p.y=q.y$ (unsatisfied). For the second case we have $p.x=q.x$ (unsatisfied) and $p.y=q.y$ (unsatisfied), and an additional constraint $p.a=q.a$ (unsatisfied). The first solution would be preferred over the second under LPB.

We also considered WSB for records. The error for $\{x:5\}=\{x:5\}$ or $\{x:5\}=\{x:6\}$ is clear enough (0 and 1 respectively). But what is the error for $\{x:5\}=\{x:5,y:6\}$ or $\{x:5\}=\{y:6\}$ or $\{x:5\}=\{x:"squid"\}$? We need to determine an error for additional or missing fields, and for comparing different types and this error needs to be such that it can be meaningfully compared with the error for values of fields that hold numeric types. Is the error for $\{x:0\}=\{x:1000000000\}$ less than for $\{x:0\}=\{y:0\}$? We could come up with a definition, but it's not clear it's that useful.

Another problem with having the solver handle records without filtering out weird cases with structural compatibility checks is that it may result in less predictability for programmers. Consider:

```
p := {x:0};
q := {y:1};
always p=q;
```

If this isn't rejected by a structural compatibility check, there are a number of plausible values that the solver could find for p and q — namely $\{x:0\}$, $\{y:1\}$, and $\{x:0,y:1\}$ — and unclear which is better. The non-determinism affects the structure of the result rather than just its value, and would prevent achieving Goal 3 in our list in Section 3.2.

Yet another complication is that we would want a minimality condition on records. Suppose we let an `always` constraint create a new record if need be:

```
always p.x = 5;
```

We want the solution to be $p=\{x:5\}$, but not e.g. $p=\{x:5, y:1000\}$.

These issues, along the lack of clear use cases for having a solver untamed by structural compatibility checks, led us to reject the alternative outlined in this subsection.

For Babelsberg/Objects, the analog of this behavior would be to allow the constraint solver generate new classes on the fly. Particularly since there would be multiple possible classes, as before, this seems overly complex and unpredictable.

A.5 Backwards Compatibility Mode for Methods

Methods called by constraints must be free of side effects. (In a practical language, benign side effects might be allowed, but we don't model this in the formal semantics.) This seems like a reasonable restriction that we shouldn't attempt to change, although we probably ought to make precise what constitutes a benign side effect.

In addition, if a method consists of more than just a single return statement, it can only be used in the forward direction in constraints. This provides a straightforward way for both the programmer and the implementor to understand how this should be handled: if the method is used in the forward direction only, it can be called in the ordinary way; if it is potentially used multi-directionally, it is exploded.

If we are trying to add constraints to an existing language with a substantial class library, this restriction does imply that many methods, even though they are side-effect-free, can only be used in the forward direction. The Babelsberg/Ruby implementation provides a “backwards compatibility mode” that relaxes the restriction in an attempt to make more of these existing methods be usable multi-directionally. (It won’t allow an arbitrary method to be successfully used in all modes in a constraint, but it increases the chances that it will work.) However, this mode is incompatible with the new semantics for Babelsberg described in this memo. We describe it here for completeness; updating it to be compatible with the new semantics (or simply removing it) is left for future work.

The Babelsberg/Ruby implementation uses two interpreter modes: standard imperative execution mode and constraint construction mode. When methods are evaluated in constraint construction mode, assignment statements are converted to two-way equality constraints. However, we’ve now decided that methods called from a constraint expression cannot themselves create other constraints, or call further methods that do so (Section 7.5), which rules out this transformation in its current form.

For the record, though, if we relaxed this restriction, we might want to support this conversion. It still doesn’t seem entirely clean, so we should probably still omit it from the formal semantics. Instead, a practical implementation of Babelsberg might elect to include a backwards compatibility mode for methods that does this transformation, as well as other transformations that increase the chances the method can be used in multi-directional constraints. In addition to converting assignments to constraints, the transformation should create fresh variables each time a variable is used on the left-hand-side of an assignment, and systematically use that new variable thereafter (until another assignment). First, here is an example that just uses straight-line code.

```
def add_and_double(x,y,z)
  sum := 0;
  sum := sum+x;
  sum := sum+y;
  sum := sum+z;
  return 2*sum;
end;
```

In backward compatibility mode this is rewritten as:

```
def add_and_double(x,y,z)
  sum_1 = 0;
  sum_2 = sum_1+x;
  sum_3 = sum_2+y;
  sum_4 = sum_3+z;
  return 2*sum_4;
end;
```

Notice that a fresh variable is introduced for each assignment, but not for the return statement. Now, if we evaluate the constraint `always 100 = add_and_double(10,15,n)`, the system should satisfy the constraint by making `n` be 25.

Conditionals, in which a variable might or might not be assigned to, can be handled by using a required equality constraint to pass through the old value for the other branch of the conditional. For example:

```

def maybe_double(x)
  ans := x;
  if x<10
    then ans := 2*ans;
  end;
  return ans;
end;

```

This becomes:

```

def maybe_double(x)
  ans_1 = x;
  if x<10
    then ans_2 = 2*ans_1;
    else ans_2 = ans_1;
  end;
  return ans_2;
end;

```

Note though that the semantics of conditional statements are unchanged — we still eagerly evaluate the test and select the appropriate branch (with no backtracking to the other branch).

If we also had a way to explode methods on demand to support recursion, we might also be able to automatically convert methods with loops to recursions that could then be used multi-directionally. Here is a version of `sum` that might be automatically generated in this fashion:

```

def sum
  helper(0,i,0,ans);
  return ans;
end;

def helper(old_i, new_i, old_ans, new_ans)
  if old_i<self.length
    then
      temp_ans = old_ans + self[old_i];
      temp_i = old_i + 1;
      helper(temp_i,new_i,temp_ans,new_ans);
    else
      new_ans = old_ans;
      new_i = old_i;
    end;
end;

```

Here, `helper` is produced automatically from the `while` by parameterizing the helper method with the old and new values of `i` and `ans`, and converting the `while` to an `if` with a recursive call. (Again, we still eagerly evaluate the test in the conditional, i.e., we evaluate the original `while` statement a definite number of times — this method won't let us try different array lengths to satisfy the constraints.)

As we've been noting, these transformations won't work in all cases. As an extreme example, if we have an `encrypt` method that takes a plaintext message `m` and a public key `k`, we are unlikely to be able to run

the method backwards to find the original message given the cyphertext. Coming up with a good set of transformations, a specification of when they will work, and whether the result covers a useful set of cases, is an open problem. Certainly they work for methods consisting just of a sequence of assignment statements and a final return statement, but this seems like not that useful a class of methods to run backwards.

For the record, here are some more straightforward examples of when the transformation fails.

Consider an absolute value method for integers:

```
def abs()
  if self>=0 then return self else return 0-self;
end;
```

Suppose we use it in this program:

```
x := 5;
y := 5;
always y=x.abs();
always x<0;
```

This works up until the final statement. However, when we try to satisfy $x < 0$, because of the semantics of the `if` statement (which eagerly evaluates the test, with no backtracking), we incorrectly conclude the constraints are unsatisfiable.

Instead, the `abs` method as written should only be used in the forwards direction; if we want one that also works backwards, we should write it as a disjunction that can be turned over to the solver (or add support for Prolog-style backtracking — see the next section).

Another case that doesn't work arises when the method has a conditional, and the two branches return different types. Even if these are both value classes, this still doesn't explode correctly using the current rules. For example, suppose that in addition to standard points, we have a class `ZeroPoint` representing a point with its `x` and `y` both 0 (which then avoids storing those fields). Also suppose that both kinds of points implement an `optimize` method that returns a new object, which is the point represented in the optimal way. For ordinary points, the method is:

```
def optimize()
  if self.x=0 && self.y==0 then return ZeroPoint() else return self;
end;
```

For `ZeroPoint`, `optimize` just returns `self`.

Now consider this code that uses the `optimize` method:

```
p := Point(5,5);
q := Point(5,5);
always q = p.optimize();
```

If we implement this by exploding the `optimize` method, the existing transformations don't work since `q` might be either a `Point` or a `ZeroPoint`. We might be able to get it to work in this case, but the transformations become more complex and probably more fragile, and it's not clear that the complexity is worth the price.

A.6 Adding New Solvers and Extending the Solver Language

In a practical Babelsberg implementation, it is useful to be able to add new solvers, if necessary extending the solver language so that new kinds of constraints can be encoded and sent to the new solvers. This doesn't seem to create any issues for the formal semantics, so we note this only in this appendix.

One useful extension of this sort would be useful to include a solver for geometric constraints, either in 2 or 3 dimensions.

Some extensions that already exist in the practical languages are a solver (and constructs in the solver language) for finite domain constraints, as in the `clpfd` library for SWI Prolog. Other additional solvers are available for strings and local propagation constraints, as solved for example by DeltaBlue. Supporting the latter required adding a construct for declaring local propagation constraints (including the propagation methods) in the source language, and also adding local propagation constraints to the solver language.

In Babelsberg/JS, when a DeltaBlue constraint is constructed, and an equality constraint is encountered with two variables of the same type on either side, when processing a constraint the interpreter doesn't descend further, but instead asserts the equality on those variables. So, for example, for two points the constraint expression does not desugar into multiple constraints on the `x` and `y` variables. This works even if the programmer does not specify a propagation function, because DeltaBlue has a default propagation method for equality (and some other relations, such as scaling.)

Finally, we may be able to have a solver language and solver to encode Prolog-like goals, which would support backtracking and Prolog-style programming within Babelsberg. A key observation here is that this would be in a separate set of constraints — we wouldn't try to change the basic Babelsberg semantics to allow backtracking with Babelsberg `if` statements.