

Declaring Constraints on Object-oriented Collections

TIM FELGENTREFF¹ ROBERT HIRSCHFELD¹ MARIA GRABER¹ ALAN BORNING²
HIDEHIKO MASUHARA³

Abstract: Logic puzzles such as Sudoku are described by a set of properties that a valid solution must have. Constraints are a useful technique to describe and solve for such properties. However, constraints are less suited to express imperative interactions in a user interface for logic puzzles, a domain that is more readily expressed in the object-oriented paradigm.

Object constraint programming provides a design to integrate constraints with dynamic, object-oriented programming languages. It allows developers to encode multi-way constraints over objects using existing, object-oriented abstractions. These constraints are automatically maintained at run-time.

In this paper we present an application of this design to logic puzzles in the Squeak/Smalltalk programming environment, as well as an extension of the design and the formal semantics of Babelsberg to allow declaring constraints using the imperative collection API provided in Squeak. We argue that our implementation facilitates creating applications that use imperative construction of user interfaces and mutable program state as well as constraint satisfaction techniques for different parts of the system. The main advantage of our approach is that it moves the burden to maintain constraints from the developer to the runtime environment, while keeping the development experience close to the pure object-oriented approach.

Keywords: Object Constraint Programming, Constraint Imperative Programming, Constraint Solving, Babelsberg

1. Introduction

Logic puzzles are declarative. Their rules declare *what* a valid solution should look like, and they can then be solved without any pre-described algorithm other than logical deduction techniques. A famous example is Sudoku. The rules of a logic puzzle describe properties that should be maintained while solving the puzzle. For example, in Sudoku, the properties are that each row, column, and box contain the numbers from 1 to 9 exactly once. The properties of a logic puzzle can be formulated as formal constraints, which a constraint solver can use to find one or more solutions or to check if a solution input by the user is valid [12].

User interface frameworks such as *Morphic* [15] are inherently imperative – the user interface consists of compositions of *Morphs* that have state and react to user input events. *Morphic* was first implemented in Self, with later implementations in Squeak [16] and JavaScript [19].

Babelsberg [5] is a design to integrate constraints into object-oriented languages in a way that allows programmers to dynamically create and satisfy constraints on objects. The design is a strict extension of the object-oriented semantics of the underlying host language. Babelsberg uses object-oriented method definitions to define constraints rather than

a constraint domain-specific language (DSL) [17], [18]. As a consequence, Babelsberg respects encapsulation and object-oriented abstractions. The design also supports solver features such as constraint priorities [2] and incremental resolving [8]. Recently, the design has been extended to allow multiple constraint solvers to cooperate to find a solution [6].

This design lends itself well to build interactive user interfaces for logic puzzles where the puzzle rules are expressed as constraints on the *Morphic* objects. In a standard imperative programming language, constraint solving and satisfaction is implemented explicitly. Using just *Morphic* in a standard imperative language, developers have to ensure that all event sources that might change the user interface resatisfy constraints or call an external constraint solver. In contrast, Babelsberg maintains constraints automatically, regardless of how the system was perturbed. This reduces the amount of knowledge the developer has to have about possible event sources for the *Morphs*. We argue that this is more in line with the encapsulation and abstraction desired in object-oriented applications.

An incomplete aspect of the original Babelsberg design was that it only allowed constraints on objects and their parts, but did not allow multi-directional solving for constraints on collections. In the context of logic puzzles the rules are usually defined on sets of objects (for example, Sudoku constraints are defined on rows, columns, and boxes.) In prior work, we experimented with an extended Squeak/Smalltalk based prototype implementation of the

¹ Hasso Plattner Institute, University of Potsdam, Germany

² University of Washington, Seattle, WA, USA

³ Department of Mathematical and Computing Sciences, Tokyo Institute of Technology

Babelsberg design — Babelsberg/S — to support operations on collections of objects [9]. In this paper, we derive a general design from this prototype implementation, as well as semantic rules to supplement the existing Babelsberg semantics [7].

Thus, the contributions of this work are:

- We describe an implementation of the Babelsberg design in Squeak/Smalltalk.
- We describe an extension to Babelsberg that let the programmer conveniently specify constraints on collections.
- We present a technique for Morpich applications to interact with constraints, using as a running example an interactive Sudoku application
- In an appendix, we present semantic rules to supplement the formal Babelsberg design to support collection predicates.

2. Object Constraint Programming in Squeak

This section describes how constraints are expressed in our Squeak implementation of Babelsberg, called Babelsberg/S. For our examples, we use the rules of a Sudoku puzzle.

```

1 constraint := [
2   (sudoku at: 1 at: 1) between: 1 and: 9
3 ] alwaysSolveWith: solver.
```

Listing 1: Defining the domain of a Sudoku cell

Listing 1 shows the constraint for defining the domain of one Sudoku cell. In general, a constraint in Babelsberg/S is specified as a block that evaluates to a boolean — if the block evaluates to `true`, the constraint is satisfied. As mentioned in Section 1, this block contains Smalltalk code, rather than code written in a separate DSL. The variable `sudoku` in Listing 1 represents the grid of cells in the interactive application from the outside scope and the method `between:and:` is a predefined predicate on Squeak numbers that just checks whether the receiver’s value is between the upper- and lower-bound arguments. To actually turn this code in into a constraint that can be handed to a solver, we send the message `alwaysSolveWith:` to the block, passing as argument an instance that implements the interface of the Babelsberg/S `ConstraintSolver` class. (It is also possible to solve the constraint with a default constraint solver, which is global inside the Squeak image, by sending `alwaysTrue`.) While the Smalltalk block can contain arbitrary Smalltalk code, asking the system to interpret it as a constraint puts the same restrictions on the expressions inside the block as for previous implementations of Babelsberg [5], [6]. These are a) an expression that is used as a constraint must evaluate to a boolean (the constraint is that it evaluate to true), b) the expression should return the same result on repeated evaluation (so that, for example, a random number generator would not qualify), and c) the expression should be free of side-effects.

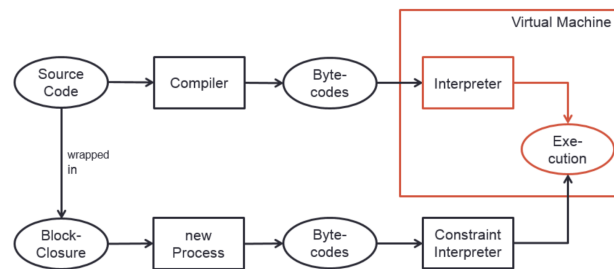


Fig. 1: The architecture of Babelsberg/S constraint construction mode

Translating Constraints in Babelsberg/S

To translate the Smalltalk expression into a form suitable for a constraint solver, the constraint block is executed in a different execution mode called *constraint construction mode* which uses symbolic execution [3], [13] to create constraint expressions from the code. The block is only evaluated in constraint construction mode when either `alwaysTrue` or `alwaysSolveWith:` are sent to it, otherwise it is just an ordinary Squeak block.

Squeak/Smalltalk includes an in-image Smalltalk interpreter that we instrumented to implement constraint construction mode. The resulting architecture is shown in Figure 1. Squeak stack frames can be reified into instances of subclasses of the `ContextPart` class. These provide methods to interpret each bytecode. This facility is used by the Squeak debugger. Babelsberg/S uses the instrumented interpreter to evaluate the constraint block. The `alwaysTrue` method creates a new Process (a Smalltalk green-thread) that is interpreted stepwise using the interface of the `ContextPart` objects. Where interpretation in constraint construction mode deviates from normal Smalltalk semantics, we use `ContextS` [10] to instrument methods whose behavior needs to change inside a constraint construction mode layer.

Consider the above constraint: the block `[(sudoku at: 1 at: 1) between: 1 and: 9]` is compiled into bytecode. A new Squeak process is created (but not scheduled) by sending the method `newProcess` to it. The process has a stack with exactly one frame (a `ContextPart` object.) That frame’s program counter is set to 0 and it contains the bytecode for the constraint block. The Babelsberg/S interpreter then steps through this frame by interpreting the bytecodes one by one, including doing method lookup and creating new frames as needed. An important consequence of this is that a variable binding that is used as receiver in a constraint block cannot be allowed to change, because then the lookup, and thus the constructed constraint, might be invalid. Thus, for Listing 1, the solver cannot simply find a collection that already satisfies the constraint and change the binding of the `sudoku` variable. Instead, it has to change the contents of the Sudoku to satisfy the constraint. This restriction does not apply to bindings that were created during constraint construction, such as return values of methods – so the solver can (and

will) change what the method `at:at:` returns when sent to `sudoku`.

The modified interpreter creates `ConstraintVariable` objects for instance variables that are accessed through accessor methods. All methods are then called on these `ConstraintVariable` objects. Operator methods such as `+`, `-`, or `<=` construct constraint expressions instead of evaluating directly. Other methods that the solver does not directly support are partially evaluated to break them down into the primitive operations. In the case of `between:and:`, for example, the constraint constructed from partially evaluating the method would be equivalent to specifying `n >= lower and: [n <= upper]` directly. By re-using existing methods, Babelsberg/S supports the object-oriented abstractions that already exist in the system. This is equivalent to the Babelsberg implementations in Ruby and JavaScript [5].

Additionally, the interpreter creates instance-specific method wrappers to intercept access to these variables. The wrappers delegate read and write access to the corresponding `ConstraintVariable`, which calls the solver as needed to keep the constraints satisfied and returns the value of the variable from the solver’s solution.

In contrast to JavaScript or Ruby, Squeak/Smalltalk does not allow instance-specific behavior directly. All methods and instance variables are declared on the class. However, wrapping accessors on the class of any encountered object would cause all instances of that class in the system to go through our wrapper, which imposes considerable performance overhead. To wrap only the encountered instances, we create anonymous subclasses of their class, and use Smalltalk’s `become:` facility to change the class of the object to the anonymous subclass. We then install our wrappers only on this instance-specific subclass.

This solution to instance-specific behavior means that there is no run-time overhead when using objects that have no constraints on them. Constrained objects are easily discovered through Smalltalk’s meta-programming interface, because their class has no name and only wraps the accessors encountered in the constraint. We encountered methods in the core system that check for the class of its arguments not using the `isKindOf:` method (which works correctly for instances of subclasses), but by directly comparing the class pointer. Although one might consider this as a bug in the method, we are working on a solution to instance-specific behavior that is completely transparent to these common uses of meta-programming.

After constraint construction has interpreted the block, the generated constraint expressions are added to a `Constraint` object, which is passed to the constraint solver. We explain the solving process in more detail in Section 3.3. If solving succeeds, the method `alwaysSolveWith:` returns the newly created constraint object. This object can then be used for reflection (e.g., to inspect which variables participate in the constraint) as well as to dynamically disable and re-enable the constraint. If solving fails, an exception is

raised, which must be handled by the programmer. In that case, the constraint is not added and the system remains unperturbed.

3. Constraints on Collections of Objects

The original Babelsberg design did not support constraints on collections directly; rather, it was proposed to use a specialized solver for collections [5]. To model an entire Sudoku puzzle, we need to assert the constraint given in Listing 1 for each cell. With the existing Babelsberg design, this would either require a solver for collections that supports domains for numbers, or alternatively, loop over the cells imperatively (Listing 2.)

```

1 (1 to: sudoku size) do: [:index |
2   [(sudoku at: index) between: 1 and: 9]
3   alwaysSolveWith: solver].

```

Listing 2: Defining the domain of all Sudoku cells with a loop

The code has two main problems, however. First, if we consider collections that can grow (or shrink), these constraints would then be incorrect — they would either have to be redacted and the loops re-executed or after adding the above constraints once any changes to the size of the collection must be prohibited. Second, many object-oriented languages including Squeak/Smalltalk come with application programming interfaces (APIs) to work with collections, and rather than iterating manually, a method such as `allDifferent`, if available on the `Collection` class, should work in a constraint, because it satisfies our restrictions on constraint expressions that they return a boolean and are free of side-effects:

```

1 [collection allDifferent] alwaysTrue.

```

The formal design of Babelsberg indeed does support such methods, but only for solving in the *forward direction*, that is, to execute them and use the result as a constant [7]. In this case, solving in the forward direction would be of little use, however, since if the result of the call to `allDifferent` is not already true, there is nothing the system can do.

Supporting collections in constraints more directly is thus useful in at least this application, and more generally in any application that deals with finite domain problems as well as representations of those problems (graphical or otherwise) that are more readily expressed using imperative code. This combination makes Sudoku applications an ideal example of the kinds of applications we want to support with this design.

There are a number of collection predicates that are commonly used in constraints and that would be useful to support. For these, we propose that implementers of Babelsberg-like languages must check their actual implementation and decide for each if they should allow them in a modified form of constraint construction mode. In this mode, rather than simply executing through complex methods involving loops, we convert any operations involving collection elements that are used as tests into constraints. Any

Table 1: Mapping from collection predicates to declarative representation

<code>anySatisfy:</code>	$\exists x \in \text{array} : f(x)$
<code>noneSatisfy:</code>	$\forall x \in \text{array} : \neg f(x)$
<code>allSatisfy:</code>	$\forall x \in \text{array} : f(x)$
<code>includes:</code>	$y. \exists x \in \text{array}. x = y$

indexing variable is treated as read-only. If the test would trigger an early return, we ignore the return and continue.

Depending on the type of recognized test, the generated constraints must then added to a conjunction or disjunction. The system determines whether to use a conjunction or disjunction if the method uses an early return optimization. If, depending on an element test, the method would early return `true`, the tests must be combined in a disjunction, since it is enough to satisfy just one test to have the method return the same. Otherwise, the elements are negated and added in to a conjunction. Thus, an implementation of `allDifferent` as in Listing 3 would be turned into a conjunction of pair-wise inequality constraints.

```

1 allDifferent
2   1 to: self length do: [:i |
3     1 to: self length do: [:j |
4       (i ~= j and: (self at: i = self at: j))
5         ifTrue: [↑ false]]]
6   ↑ true

```

Listing 3: A possible implementation of `allDifferent`

The code for `allDifferent` would be expanded into a conjunction of constraints, because the early return is `false`. The constraints in the conjunction would be for the tests to that early return, pair-wise constraining `self at: i = self at: j` to be false (with the values of `i` and `j` fixed for each constraint). In addition, supposing the `length` method represents a field access, this field is also used in the constraint, and thus the system can track any change to this field to trigger regenerating the constraints on the collection.

Some common predicates available on collections in Squeak/Smalltalk are translated as per Table 1. (Note that both the test if some element satisfies a particular predicate as well as the test for membership (the latter being a special case of the former) require a disjunction. Without any other constraints, and since our design does not use Prolog-style backtracking, they would probably always be satisfied by setting the first (or last, depending on the concrete implementation) element of the array.) Even though this list contains only a few predicates, in practice many languages come only with a small set of primitive collection types that are supported at a language level. Languages with a rich collection library such as Common Lisp or Squeak/Smalltalk [11] are built around a small number of types and primitive operations to access and store indexed elements in an object. Thus, implementing the special support needed to support these basic predicates enables their use in a variety of contexts, including methods that are built on top of these predicates.

3.1 Constraints on User-Defined Methods

We do not intend for language implementers to support every possible method that a collection may have in a practical implementation, in particular if that collection may be extended with user-defined methods. As an example, consider an iterative `sum` method as in Listing 4.

```

1 sum
2   | answer |
3   answer := 0.
4   1 to: self length do: [:i |
5     answer := answer + self at: i].
6   ↑ answer

```

Listing 4: A possible implementation of `sum`

We can of course use this method in the forward direction in a constraint:

```

1 a := Array new: 2.
2 a at: 1 put: 10.
3 a at: 2 put: 20.
4 s := 30.
5 [s = a sum] alwaysTrue.
6 a at: 1 put: 100.

```

After the constraint is executed, `s` is 30 (since the constraint is already satisfied); then after setting the first element of `a` to 100, `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:

```

1 a := Array new: 2.
2 a at: 1 put: 10.
3 a at: 2 put: 20.
4 [50 = a sum] alwaysTrue.

```

We have found a design pattern for user code that works well in these situations that *can* provide something that works both forward and backward. Rather than using a method that returns the calculated sum of the elements, we eagerly update the sum as the array changes in a variable. This can be done by writing an ordinary method that sets up a recursive network of addition constraints over the array elements. The `sum` becomes an instance variable of our collection, and the implementation of that collection must take care to correctly initialize the constraint network when it is created or an element is added or removed. An example of such an initialization is given in Listing 5. The advantage for code using the collection's `sum` in further constraints is that it is used simply as a variable — constraints on it can work both ways, and it can even be assigned and the array changes to satisfy the constraints.

```

1 initializeSum
2   self.length = 0
3   ifTrue: [[self sum = 0] alwaysTrue]
4   ifFalse: [[self sum = (self at: 1) +
5             self allButLast sum] alwaysTrue].

```

Listing 5: A possible initialization for a constrainable `sum`

3.2 Implementing Constraints on Collections in Babelsberg/S

Babelsberg/S implements a prototype of our scheme to support constraints on object-oriented collections. As a result, the domain constraint of a Sudoku puzzle can be expressed through sending the collection predicate `allSatisfy:` to `sudoku` (Listing 6).

```

1 [sudoku allSatisfy: [:cell | cell between: 1 and: 9]]
2   alwaysSolveWith: solver.
```

Listing 6: Defining the domain of all Sudoku cells with the Collection API

The extension to support collections directly in Babelsberg/S leverages the fact that Smalltalk comes with only one fixed-size pointer array type, upon which the Smalltalk collections library builds. This type provides three methods implemented in primitives for all low-level access: `at:`, `at:put:`, and `replaceFrom:to:with:startingAt:`.

Babelsberg/S subclasses the basic `Array` class and overrides the three low-level access methods to intercept any modifications to the array. In addition, it overrides the `copyFrom:to:` method, which is regularly used in Squeak to access sub-sequences of an array.

In constraint construction mode, any array that is visited in the dynamic extent of the execution is transparently replaced by the Babelsberg/S subclass. Besides the overridden methods, this subclass is a completely transparent proxy. Each element in the array that participates in the constraint is wrapped in a `ConstraintVariable`. The predicates of the collection API are straightforward to support. The predicates `anySatisfy:`, `noneSatisfy:`, and `allSatisfy:` are mapped as per Table 1. Note that for the first relation, a disjunction over all elements must be created. For solvers that do not support disjunctions, Babelsberg/S forces the first element to satisfy the block. This prevents the system from finding solutions in many cases. To find additional solutions with solvers without disjunctions requires backtracking in the case of unsatisfiable constraints. This is not implemented yet, but is left for future work.

In general, any predicate method available on collections can be used in constraints, as per our design for supporting user defined methods as collection predicates. For example, predicate methods such as `allDifferent:` can be mapped to pair-wise inequalities by simply interpreting their implementation in constraint construction mode. Other methods that are useful in constraints reduce all elements of a collection and then express properties over those reductions. Reduction methods include the `sum` method mentioned above, as well as the `count:` method that returns the number of elements that satisfy a particular condition.

The constraints created with these methods are reconstructed when the elements in the array change, but since the size of arrays is fixed, the length of the linear expressions is bounded, so in the Babelsberg/S implementation, we only initialize intermediate variables once, as the underlying array cannot grow or shrink. Since we only need to initialize

them once, this can be taken care of by the framework.

Predicates over expressions are useful to state constraints on a collection as a whole, rather than on each of its elements. We have used this, for example, in our implementation of the Outside-Sum-Sudoku. Here, all elements in a collection must sum to the number outside the Sudoku. When one element changes, the others must change, too, to ensure the total sum does not.

We have found few use-cases for the most general collection-methods `do:`, `collect:`, and `select:` that could not be expressed using more specific methods. These methods create new collections from existing ones. What the developer means when using them and how to translate that meaning to the solver is less clear in the general case, and we do not support them for now. We have found that uses of `select:`, `collect:`, and `detect:` in predicate expressions can usually be replaced by the direct predicate methods. We have found that the iteration method `do:` is usually just used to express constraints on each element, and can usually be pulled outside of the constraint block. We might lift this restriction in the future if we find a significant number of uses of these methods in constraints that are much more expressive than their direct predicate counterparts. Until then, and for simplicity in the implementation, we do not allow these methods in constraints.

3.3 Maintaining Constraints in Babelsberg/S

Once asserted, constraints need to continue to be satisfied until they are disabled, all objects they apply to are garbage collected, or the program stops. To ensure this, the Babelsberg design follows the *perturbation* model established by the Kaleidoscope constraint-imperative language [14]. This model is similar to reactive systems in that changes to one part of the system propagate to other parts. In reactive systems, these changes are made by sampling a continuous process or through discrete events. In Babelsberg, the changes are the concrete event of assigning a new value to a variable that participates in a constraint. These changes then potentially propagate to other variables to keep constraints satisfied.

Each variable that participates in a constraint implicitly reacts to programmatic changes to its value by calling one or more solvers to re-satisfy the variable's constraints. Our wrappers around accessors intercept changes to variables that were used in a constraint and call `suggestValue:` on their associated `ConstraintVariable`. This adds a temporary equality constraint for the new value to the underlying constraint solver. The solver tries to solve all constraints. If the constraints are satisfiable, the new value is assigned. As a side-effect, other variables might change to satisfy constraints. If the solver cannot find a solution, a runtime error is generated and the new value is ignored.

In our Sudoku example, if a new value is assigned to a cell, the Sudoku constraints are solved in the background. If the solver finds a solution, the cell changes its value and the rest of the puzzle is adjusted to keep the Sudoku solveable.

4		9		1	6			
6			3	8	9		4	
	7			4	5			9
				6		5	2	1
		4		7		6		
1	9			5				
9			4	2	1		7	3
	3		6	9	8			
				3				6
Give a hint								

Fig. 2: The Morphic UI of the Sudoku puzzle. Some numbers (in black) are given initially. Each free cell allows the user to input a single number (in blue). The system can also generate hints (in red).

That is possible because the Sudoku application interacts with the underlying constraints.

Babelsberg/S can accommodate a variety of constraint solvers. Currently, it supports the Squeak implementation of Cassowary [1], and Z3 [4] through an IPC interface.

4. Evaluation

The constraints in Sudoku are easy to state, but not always easy to satisfy. A correct solution must assign each cell a number between 1 and 9 inclusively, while at the same time ensuring the no number occurs twice in a row, a column, or a block of 3 by 3 cells. We argue that logic puzzles such as Sudoku are good examples for interactive constraint applications. The user interface (written in Morphic) is shown in Figure 2.

Listing 7 shows the constraints necessary to solve this Sudoku puzzle. These constraints use the Z3 constraint solver. Line 1 ensures that the user cannot change the numbers that were given initially. In some solvers, such as Cassowary, stay constraints can be used to express that the solver may not change a given variable, or to only change it if the constraints cannot be satisfied otherwise. Stay constraints are currently not supported in Z3, but will be in future versions. Currently, the method `addConstraintsForAllGivenNumbers` iterates over cells and creates a constraint that each cell that already has a value is always equal to just that value. Lines 3–4 assert the constraint that all cells must contain numbers between 1 and 9. Finally, lines 6–10 ensure that no row, column, or 3×3 box of cells can have duplicate numbers.

Note that the Squeak collection API does not contain a method `allDifferent`. Babelsberg/S adds this predicate for convenience. It is a normal object-oriented method in the `Collection` class that iterates over all elements in the collection and tests them for pairwise inequality. In ordi-

nary code, this is just a test – the constraint interpreter, however, creates an inequality constraint expression for each comparison, exploding the `allDifferent` method into multiple constraints that the solver can understand. This means also, that subclasses can override the method and any different behavior will be reflected in the created constraints.

Note also that the normal accessor methods for rows and columns from Squeak `Matrix` objects are used, too. The Sudoku grid is just a subclass of `Matrix` that, besides a method to assert constraints on the given numbers, adds the `atBox:` accessor method to access each of the 9 boxes of size 3×3.

```

1  sudoku addConstraintsForAllGivenNumbers.
2
3  [sudoku allSatisfy: [:cell |
4    cell between: 1 and: 9]] alwaysSolveWith: solver.
5
6  (1 to: sudoku rowCount) do: [:index |
7    [(sudoku atRow: index) allDifferent &
8     (sudoku atColumn: index) allDifferent &
9     (sudoku atBox: index) allDifferent]
10   alwaysSolveWith: solver].

```

Listing 7: All Constraints of a Sudoku Puzzle

As can be seen from Listing 7, the amount of code necessary for specifying all properties of a Sudoku puzzle is very small. With these, a solver can solve an arbitrary given Sudoku puzzle. The constraints are completely decoupled from the specific Sudoku puzzles and their given numbers.

In Babelsberg, constraints can be constructed, enabled and disabled at run-time, and, because they work correctly with method polymorphism, it is possible to subclass a logic puzzle to construct another by adding or removing constraints only. As an example, we have created Sudoku puzzle subclasses for *Diagonal Sudokus* and *Outside-Sum Sudokus*. In the former, the numbers of the two main diagonals have to be all different, and in the latter, the first three numbers in a row or a column must add up to a specific sum.

For a Diagonal Sudoku, provided there are accessor methods for the two diagonals, the method to create constraints is shown in Listing 8.

```

1  super createConstraints. "from normal Sudoku"
2  [(self diagonalFromTopLeft allDifferent)
3   and: [self diagonalFromTopRight allDifferent]]
4   alwaysSolveWith: solver.

```

Listing 8: The Diagonal Sudoku

With object-constraint programming (OCP), it does not matter in which way a constraint variable or a constraint changes. The constraint satisfaction automatically works on each disturbance of the system. Currently, the values of cells only change when the user enters a new value into the morph that represents a cell. If that value is not a number between 1 and 9, or the Sudoku cannot be solved by adding this value, the solver rejects the input. However, the constraints encode no source for the change, so it does not matter if the change actually occurred through keyboard input. The Sudoku could also be calculated entirely by the computer, or the game could allow remote users to send values over the network. The constraints thus provide flexibility,




Fig. 3: Squeak debugger with constraints

because the developer does not need to know all events that might change the puzzle.

5. Conclusions

We have argued that OCP facilitates reactive systems in which dependencies between objects can be declared as constraints. It modularizes the relationship between objects and decouples constraint satisfaction from the application. Constraints can be dynamically added and removed, and are maintained automatically. This makes them useful for writing interactive applications. As an example, we implemented applications for specifying and solving different variants of Sudoku with constraints with a graphical user interface. The user can change the values of the constraint variables interactively without breaking the properties of the Sudoku. The application reacts on the user input by resolving the underlying constraints.

There are two major directions for future work. Regarding the implementation, we plan to implement an alternative solution to provide instance-specific wrappers. This will improve the compatibility of constrained objects with existing Smalltalk code. We also plan to support more features of the Babelsberg design as found in its JavaScript and Ruby implementations, such as incremental resolving, local propagation, and identity constraints.

Furthermore, we plan to leverage the Smalltalk meta-programming facilities to explore how to aid developers in debugging and understanding constraints. If incorrect constraints are generated, why? If the solver cannot find a solution or is slow, what can be done? These are still open questions for Babelsberg, because constraints cannot be easily debugged.

Figure 3 shows how we extended the Squeak debugger to support stepping into constraints. Our debugger has an additional pane on the top right ①, and a special inspector on the right hand side ②. The debugger works as it normally would when running imperative code, but upon entering constraint construction mode, the debugger additionally tracks the constraints as they are created. In the example, we assert that `pt1` and `pt2` should be equal ③. From just looking at the expression we cannot tell how many constraints would be created. We could infer from the implementation if the `=` method for 2d points that we will create

two constraints, one for each pair of dimensions on those two points, but a debugger allows us to observe this fact and see the equations that have been translated for the Cassowary solver in the top right pane. On the right hand side, we can see the details of the first constraint and for example change its strength or the generated expression to see how that changes the program behavior. Additionally, we can step into the procedure that assigns updated values from the constraint solver to the program variables and thus see the global effects of a constraint. This is particularly useful to understand which solution a solver chooses for a particular constraints and how many variables are changed in which way. We plan to extend this prototype into a debugger that is useful to answer different questions that arise when developing with constraints.

Despite these avenues for future work, we think that Babelsberg/S is already a useful implementation of object-constraint programming and we plan to include it in a future release of the R/Squeak distribution, a Squeak distribution that includes research projects considered useful for general purpose development*¹.

References

- [1] Badros, G. J., Borning, A. and Stuckey, P. J.: The Cassowary Linear Arithmetic Constraint Solving Algorithm, *ACM Transactions on Computer-Human Interaction (TOCHI)*, Vol. 8, No. 4, pp. 267–306 (online), DOI: 10.1145/504704.504705 (2001).
- [2] Borning, A., Freeman-Benson, B. and Wilson, M.: Constraint Hierarchies, *Lisp and Symbolic Computation*, Vol. 5, No. 3, pp. 223–270 (online), DOI: 10.1007/bf01807506 (1992).
- [3] Clarke, L. A.: A system to generate test data and symbolically execute programs, *IEEE Transactions on Software Engineering*, Vol. 2, No. 3, pp. 215–222 (online), DOI: 10.1109/tse.1976.233817 (1976).
- [4] de Moura, L. and Bjørner, N.: Z3: An Efficient SMT Solver, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, Springer, pp. 337–340 (online), DOI: 10.1007/978-3-540-78800-3.24 (2008).
- [5] Felgentreff, T., Borning, A. and Hirschfeld, R.: Specifying and Solving Constraints on Object Behavior, *Journal of Object Technology*, Vol. 13, No. 4, pp. 1–38 (online), DOI: 10.5381/jot.2014.13.4.a1 (2014).
- [6] Felgentreff, T., Borning, A., Hirschfeld, R., Lincke, J., Ohshima, Y., Freudenberg, B. and Krahn, R.: Babelsberg/JS, *Proceedings of the European Conference on Object-oriented Programming (ECOOP)*, Springer, pp. 411–436 (online), DOI: 10.1007/978-3-662-44202-9.17 (2014).
- [7] Felgentreff, T., Millstein, T. D., Borning, A. and Hirschfeld, R.: Checks and balances: constraint solving without surprises in object-constraint programming languages, *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, ACM, pp. 767–782 (online), DOI: 10.1145/2814270.2814311 (2015).
- [8] Freeman-Benson, B. N., Maloney, J. and Borning, A.: An Incremental Constraint Solver, *Communications of the ACM*, Vol. 33, No. 1, pp. 54–63 (online), DOI: 10.1145/76372.77531 (1990).
- [9] Graber, M., Felgentreff, T., Hirschfeld, R. and Borning, A.: Solving Interactive Logic Puzzles With Object-Constraints — An Experience Report Using Babelsberg/S for Squeak/Smalltalk, *Workshop on Reactive and Event-based Languages & Systems (REBLS)*, pp. 1:1–1:5 (2014).
- [10] Hirschfeld, R., Costanza, P. and Haupt, M.: An Introduction to Context-Oriented Programming with ContextS, *Generative and Transformational Techniques in Software Engineer-*

*1 <https://www.hpi.uni-potsdam.de/swa/trac/>
↪ SqueakCommunityProjects

- ing II, Springer, pp. 396–407 (online), DOI: 10.1007/978-3-540-88643-3.9 (2008).
- [11] Ingalls, D., Kaehler, T., Maloney, J., Wallace, S. and Kay, A.: Back to the Future: The Story of Squeak, a Practical Smalltalk Written in Itself, *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, ACM, pp. 318–326 (online), DOI: 10.1145/263698.263754 (1997).
- [12] Ist, I. L., Lynce, I. and Ouaknine, J.: Sudoku as a SAT Problem, *Proceedings of the International Symposium on Artificial Intelligence and Mathematics (AIMATH)*, Springer, pp. 1–9 (online), DOI: 10.1.1.331.458 (2006).
- [13] King, J. C.: Symbolic Execution and Program Testing, *Communications of the ACM*, Vol. 19, No. 7, pp. 385–394 (online), DOI: 10.1145/360248.360252 (1976).
- [14] Lopez, G., Freeman-Benson, B. and Borning, A.: Kaleidoscope: A Constraint Imperative Programming Language, *Constraint Programming*, Springer, pp. 313–329 (online), DOI: 10.1007/978-3-642-85983-0_12 (1994).
- [15] Maloney, J.: *Morphic: The Self User Interface Framework*, 4th edition (1995).
- [16] Maloney, J.: *An Introduction to Morphic: The Squeak User Interface Framework* (2001).
- [17] Milicevic, A., Rayside, D., Yessenov, K. and Jackson, D.: Unifying Execution of Imperative and Declarative Code, *Proceedings of the International Conference on Software Engineering (ICSE)*, ACM, pp. 511–520 (online), DOI: 10.1145/1985793.1985863 (2011).
- [18] Sadun, E.: *iOS Auto Layout Demystified*, Addison-Wesley (2013).
- [19] Taivalsaari, A., Mikkonen, T., Ingalls, D. and Palacz, K.: Web Browser As an Application Platform: The Lively Kernel Experience, Technical report, Sun Microsystems, Inc. (2008).

Appendix

A.1 Extension to the Formal Semantics of Babelsberg

The semantic rules presented here are an extension to the semantics of Babelsberg/Objects, presented in the companion technical report to [7]. This appendix should be read as an additional chapter after that companion report.

The syntax is augmented to include an element \textcircled{P} , which ranges over the core predicates on collections such as `allSatisfy:`, `anySatisfy:`, `includes:` and so on. For languages which support re-definition of the methods that come with the language, we assume that the element matches only the original implementations, not user-defined re-definitions. Furthermore, we augment the syntax to also support accessing records using expressions.

L-Value	$L ::= x \mid e.1 \mid e[e]$
Label	$l ::= \text{record label names} \mid \textcircled{P}$

Table A-1 gives an overview of the additional judgments used in this extension of the semantics. We add two opaque helper judgments. The first converts constants to label names, and the second checks if a constant value refers to an array class type. Both are defined in terms of the host language API. Note that for languages the support re-definition of core classes, the second judgment will return `false` if such re-definition has taken place.

Besides those additions, we only add the extended evaluation rules for dynamic field access and the special constraint construction mode (*ccm*) for collection APIs. Note that we

do not add a typing rule for dynamic field access — during inlining, such access are turned into ordinary field accesses, and their expressions are required to stay equal to the current value.

$$\begin{array}{c}
 \frac{\langle \mathbb{E} \mid \mathbb{S} \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid e_l \rangle \Downarrow \langle \mathbb{E}' \mid \mathbb{H}' \mid \mathbb{C}' \mid \mathbb{I}' \mid c \rangle \quad \text{asLabel}(c) = 1}{\langle \mathbb{E}' \mid \mathbb{S}' \mid \mathbb{H}' \mid \mathbb{C}' \mid \mathbb{I}' \mid e.1 \rangle \Downarrow \langle \mathbb{E}'' \mid \mathbb{H}'' \mid \mathbb{C}'' \mid \mathbb{I}'' \mid v \rangle} \\
 \text{(E-EXPFIELD)} \\
 \\
 \frac{\langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e_l \rangle \rightsquigarrow \langle \mathbb{E}', e_{C_l}, e'_l \rangle}{\langle \mathbb{E}' \mid \mathbb{S} \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid e_l \rangle \Downarrow \langle \mathbb{E}'' \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid c \rangle \quad \text{asLabel}(c) = 1} \\
 \frac{\langle \mathbb{E}'', \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.1 \rangle \rightsquigarrow \langle \mathbb{E}''', e_{C'}, e'_l \rangle}{\langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e[e_l] \rangle \rightsquigarrow \langle \mathbb{E}''', e_{C'} \wedge e_C \wedge e'_l = c, e' \rangle} \\
 \text{(I-EXPFIELD)}
 \end{array}$$

We extend the inlining judgment to also work for statements. This is used in the inlining judgment to translate calls to the well-known collection predicates. These predicates will not match the previous I-MULTIWAYCALL rule, because their implementations have more than a single return expression, so that rule is unchanged. Because we allow a limited subset of statements, including assignment to locals in inlining collection predicates, the inlining rule including statements also returns an updated scope. In addition, the constraint expressions that are returned by the inlining rule for statements are split into groups for conjunctions and disjunctions — this is required to track, based on the early returns that are encountered, whether a set of inlined expressions all need to be satisfied or if just one needs to be satisfied.

We define a helper judgments to inline collection predicates:

$$\begin{array}{c}
 \langle \mathbb{E}' \mid \mathbb{S} \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid e_0 \rangle \Downarrow \langle \mathbb{E}'' \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid v \rangle \\
 \mathbb{E}; \mathbb{H} \vdash v : \mathbb{T} \quad \text{isBasicCollection}(\mathbb{T}) = \text{true} \\
 \langle \mathbb{E}_0 \mid \mathbb{S} \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid e_1 \rangle \Downarrow \langle \mathbb{E}_1 \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid v_1 \rangle \\
 \dots \\
 \langle \mathbb{E}_{n-1} \mid \mathbb{S} \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid e_n \rangle \Downarrow \langle \mathbb{E}_n \mid \mathbb{H} \mid \mathbb{C} \mid \mathbb{I} \mid v_n \rangle \\
 e_C = (e=v \wedge e_1=v_1 \wedge \dots \wedge e_n=v_n) \\
 \text{lookup}(v, l) = (x_1 \dots x_n, s; \text{return } c) \\
 \text{enter}(\mathbb{E}_n, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, v, x_1 \dots x_n, e_1 \dots e_n) = (\mathbb{E}', \mathbb{S}_m, \mathbb{H}, \mathbb{C}, \mathbb{I}, l) \\
 \hline
 \text{preparePredicate}(\mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.1(e_1, \dots, e_n)) \\
 = (\mathbb{E}', \mathbb{S}_m, s; \text{return } c, e_C) \\
 \text{(PREPAREPREDICATE)}
 \end{array}$$

This helper rule sets up the required equalities for all the arguments and the receiver, and is essentially the same as I-CALL. As an addition, it limits any inlining to collection predicates that return a constant as a final statement. In the two rules that follow, this constant is further limited to be either `true` or `false`.

$$\begin{array}{c}
 \text{preparePredicate}(\mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.\textcircled{P}(e_1, \dots, e_n)) \\
 = (\mathbb{E}', \mathbb{S}', s; \text{return } c, e_0) \\
 \text{c} = \text{true} \quad \langle \mathbb{E}', \mathbb{S}', \mathbb{H}, \mathbb{C}, \mathbb{I}, s \rangle \rightsquigarrow \langle \mathbb{E}'', \mathbb{S}', e_1, e_C, e_D \rangle \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.\textcircled{P}(e_1, \dots, e_n) \rangle \rightsquigarrow \langle \mathbb{E}'', e_0 \wedge e_1, e_C \rangle \\
 \text{(I-POSITIVEPREDICATE)}
 \end{array}$$

Table A.1: Judgments and Intuitions of Additional and Changed Semantic Rules

Opaque Judgments

$asLabel(c) = 1$ Constant c converted into a label yields 1
 $isBasicCollection(T) = c$ When type T corresponds to a known basic collection type that is supported in constraints with predicates, c is **true**.

Constraint Solving

$\langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}', e_0, e' \rangle$
 Inlining expression e in \mathbb{S} is equivalent to e' in \mathbb{E} if e_C evaluates to true.
 $\langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s \rangle \rightsquigarrow \langle \mathbb{E}', \mathbb{S}', e_0, e_c, e_d \rangle$
 Inlining statement s is equivalent to solving conjunction of constraint expressions e_C and the disjunction of constraint expressions e_D if e_0 evaluates to true. This inlining step returns an updated environment \mathbb{E}' and scope \mathbb{S}' .

Helper Rule

$preparePredicate(\mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.l(e_1, \dots, e_n)) = (\mathbb{E}', \mathbb{S}', s; \text{return } c, e_C)$
 Preparing the method call $e.l(e_1, \dots, e_n)$ for inlining returns and updated environment \mathbb{E}' , the fresh method scope \mathbb{S}' , the method body $s; \text{return } c$, and is valid if e_C evaluates to true.

$$\begin{array}{c}
 preparePredicate(\mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.\textcircled{P}(e_1, \dots, e_n)) \\
 = (\mathbb{E}', \mathbb{S}', s; \text{return } c, e_0) \\
 \hline
 c = \text{false} \quad \langle \mathbb{E}', \mathbb{S}', \mathbb{H}, \mathbb{C}, \mathbb{I}, s \rangle \rightsquigarrow \langle \mathbb{E}'', \mathbb{S}'', e_1, e_C, e_D \rangle \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e.\textcircled{P}(e_1, \dots, e_n) \rangle \rightsquigarrow \langle \mathbb{E}'', e_0 \wedge e_1, e_C \wedge e_D \rangle \\
 \text{(I-NEGATIVEPREDICATE)} \\
 \\
 \begin{array}{c}
 \mathbb{S}(x) = x_g \quad \mathbb{E}(x'_g) \uparrow \\
 \langle \mathbb{E} | \mathbb{S} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H} | \mathbb{C} | \mathbb{I} | v \rangle \\
 \langle \mathbb{E}', \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}'', e_0, e' \rangle \\
 \mathbb{S}' = \mathbb{S} \setminus \{x, x_g\} \quad \mathbb{S}'' = \mathbb{S}' \cup \{(x, x'_g)\} \\
 \mathbb{E}''' = \mathbb{E}'' \cup \{(x'_g, v)\} \\
 e_c = (e_0 \wedge e' = v \wedge x'_g = v) \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, x := e \rangle \rightsquigarrow \langle \mathbb{E}''', \mathbb{S}'', e_c, \text{true}, \text{false} \rangle \\
 \text{(I-ASGNLOCAL)}
 \end{array}
 \end{array}$$

We use two separate rules for inlining through collection predicates that return **true** or **false** as their final statement. For methods that return **true** any disjunction, which would be created by an early return true, does not have to be fulfilled, as even without the early return the method would return true. Conversely, when the method returns **false**, fulfilling any conjunction will not suffice, because that would simply prevent an early return **false**, but not the final return statement.

Since we now allow inlining through a limited subset of statements, we add inlining rules for those. Note that these rules can only come into play through an I-*PREDICATE. Furthermore, all rules not supplied here still lead to a failure to evaluate an I-*PREDICATE rule, and fall back to the previous I-CALL rule to set up a one-way constraint on the result of the call.

$$\begin{array}{c}
 \mathbb{S}(x) \uparrow \quad \mathbb{E}(x_g) \uparrow \\
 \langle \mathbb{E} | \mathbb{S} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H} | \mathbb{C} | \mathbb{I} | v \rangle \\
 \langle \mathbb{E}', \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}'', e_0, e' \rangle \\
 \mathbb{S}' = \mathbb{S} \cup \{(x, x_g)\} \quad \mathbb{E}''' = \mathbb{E}'' \cup \{(x_g, v)\} \\
 e_c = (e_0 \wedge e' = v \wedge x_g = v) \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, x := e \rangle \rightsquigarrow \langle \mathbb{E}''', \mathbb{S}', e_c, \text{true}, \text{false} \rangle \\
 \text{(I-ASGNNEWLOCAL)}
 \end{array}$$

Assignments are only permitted to local variables. Since we can only start the statement inlining rules from a collection predicate \textcircled{P} , we start with a fresh scope and any local variable must be newly created first. In this case, assignment is turned into a required equality between the fresh variable name and the initial value. Note that we are using the expression judgment to evaluate the right-hand side, but we disallow any changes to the heap or the constraint stores.

Since we do not allow creating additional constraints even in this extended inlining mode, there is no need to solve constraints when we re-assign to a local variable. Furthermore, since re-assignments are needed for looping over collection indices, and these indices are also used to then access the collection, we create a fresh global name for every re-assigned variable. This way, every re-assignment turns into a new variable for the solver.

$$\begin{array}{c}
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, \text{skip} \rangle \rightsquigarrow \langle \mathbb{E}, \mathbb{S}, \text{true}, \text{true}, \text{true} \rangle \text{(I-SKIP)} \\
 \\
 \begin{array}{c}
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s_1 \rangle \rightsquigarrow \langle \mathbb{E}', \mathbb{S}', e_1, e_{C1}, e_{D1} \rangle \\
 \langle \mathbb{E}', \mathbb{S}', \mathbb{H}, \mathbb{C}, \mathbb{I}, s_2 \rangle \rightsquigarrow \langle \mathbb{E}'', \mathbb{S}'', e_2, e_{C2}, e_{D2} \rangle \\
 e_{C3} = e_{C1} \wedge e_{C2} \quad e_{D3} = e_{D1} \vee e_{D2} \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s_1 ; s_2 \rangle \rightsquigarrow \langle \mathbb{E}'', \mathbb{S}'', e_1 \wedge e_2, e_{C3}, e_{D3} \rangle \\
 \text{(I-SEQ)}
 \end{array}
 \end{array}$$

The skip and sequence rules are straightforward. The conjunction and disjunction expressions from the sequences are connected appropriately.

$$\begin{array}{c}
 s = \text{if } e \text{ then return true else } s_1 \\
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}', e_C, e' \rangle \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s \rangle \rightsquigarrow \langle \mathbb{E}, \mathbb{S}, e_C, \text{true}, e' \rangle \\
 \text{(I-IFTHENRETURNTRUE)}
 \end{array}$$

$$\begin{array}{c}
 s = \text{if } e \text{ then return false else } s_1 \\
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}', e_C, e' \rangle \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s \rangle \rightsquigarrow \langle \mathbb{E}, \mathbb{S}, e_C, e' = \text{false}, \text{false} \rangle \\
 \text{(I-IFTHENRETURNFALSE)}
 \end{array}$$

We only support if-clauses used as early returns in this extended inlining mode. As described in Section 3, if the

early return would return `true`, the inlined conditional is used in a disjunction, otherwise it is used in a conjunction.

$$\begin{array}{c}
 s_0 = \text{while } e \text{ do } s \\
 \langle \mathbb{E} | \mathbb{S} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H} | \mathbb{C} | \mathbb{I} | \text{true} \rangle \\
 \langle \mathbb{E}', \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}'', e_0, e' \rangle \\
 \langle \mathbb{E}'', \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s \rangle \rightsquigarrow \langle \mathbb{E}''', \mathbb{S}', e_1, e_{C_1}, e_{D_1} \rangle \\
 \langle \mathbb{E}''', \mathbb{S}', \mathbb{H}, \mathbb{C}, \mathbb{I}, s_0 \rangle \rightsquigarrow \langle \mathbb{E}''''', \mathbb{S}'', e_r, e_{C_r}, e_{D_r} \rangle \\
 e' = e_0 \wedge e' \wedge e_1 \wedge e_r \quad e_C = e_{C_0} \wedge e_{C_r} \\
 e_D = e_{D_0} \vee e_{D_r} \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s_0 \rangle \rightsquigarrow \langle \mathbb{E}''''', \mathbb{S}'', e', e_C, e_D \rangle \\
 \text{(I-WHILED0)}
 \end{array}$$

$$\begin{array}{c}
 s_0 = \text{while } e \text{ do } s \\
 \langle \mathbb{E} | \mathbb{S} | \mathbb{H} | \mathbb{C} | \mathbb{I} | e \rangle \Downarrow \langle \mathbb{E}' | \mathbb{H} | \mathbb{C} | \mathbb{I} | \text{false} \rangle \\
 \langle \mathbb{E}', \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, e \rangle \rightsquigarrow \langle \mathbb{E}'', e_0, e' \rangle \\
 \hline
 \langle \mathbb{E}, \mathbb{S}, \mathbb{H}, \mathbb{C}, \mathbb{I}, s_0 \rangle \rightsquigarrow \langle \mathbb{E}'', \mathbb{S}, e_0 \wedge e' = \text{false}, \text{true}, \text{false} \rangle \\
 \text{(I-WHILESKIP)}
 \end{array}$$

Finally, the while construct is now supported during inlining. Note that the loop condition is inlined and required to stay at its value. This prevents the solver from being able to change the loop condition to, for example, satisfy the collection predicate only on a subset of the collection.

There is an issue with these rules: they may generate constraints that are too strong. Consider the following method:

```

1 def some_or_none()
2   i := 0;
3   while i < self.length do (
4     if self[i] > 10 then return true;
5     if self[i] < 0 then return false;
6     i := i + 1
7   );
8   return true
9 end
10 always ary.some_or_none()
    
```

The constraint ensures that at least one element in the array is larger than ten, or else all elements are negative. Here, the constraint would be satisfied if:

$$\begin{array}{c}
 \exists (x, i) \in \text{ary}. (x > 10 \wedge (\forall (y, j) \in \text{ary}. \neg (y < 0) \vee j > i)) \\
 \vee \\
 \forall (x, i) \in \text{ary}. \neg (x < 0)
 \end{array}$$

But the I-POSITIVEPREDICATE rule would always require the conjunction to be satisfied, so the solver would have to solve this stronger constraints instead:

$$\forall (x, i) \in \text{ary}. \neg (x < 0)$$

We have decided to avoid additional complexity in the rules to support generating the proper constraints in these cases. The code above could easily be rewritten to use two methods which each test one property, and then use these in a disjunction. Since the set of supported collection predicates \textcircled{P} is defined as part of the language, such methods may simply not be included in that set.