

Modelling Destructive Assignments

Kathryn Francis^{1,2}, Jorge Navas², and Peter J. Stuckey^{1,2}

¹ National ICT Australia, Victoria Research Laboratory

² The University of Melbourne, Victoria 3010, Australia

Abstract. Translating procedural object oriented code into constraints is required for many processes that reason about the execution of this code. The most obvious is for *symbolic execution* of the code, where the code is executed without necessarily knowing the concrete values. In this paper, we discuss translations from procedural object oriented code to constraints in the context of solving optimisation problems defined via simulation. A key difficulty arising in the translation is the modelling of state changes. We introduce a new technique for modelling destructive assignments that outperforms previous approaches. Our results show that the optimisation models generated by our technique can be as efficient as equivalent hand written models.

1 Introduction

Symbolic reasoning has been the crux of many software applications such as verifiers, test-case generation tools, and bug finders since the seminal papers of Floyd and Hoare [6, 8] in program verification and King [9] in symbolic execution for testing. Common to these applications is their translation of the program or some abstraction of it into equivalent *constraints* which are then fed into a *constraint solver* to be checked for (un)satisfiability.

The principle challenge for this translation is effective handling of destructive state changes. These both influence and depend on the flow of control, making it necessary to reason disjunctively across possible execution paths. In object oriented languages with field assignments, the disjunctive nature of the problem is further compounded by potential aliasing between object variables.

In this paper we introduce a new, *demand-driven* technique for modelling destructive assignments, designed specifically to be effective for the difficult case of field assignments. The key idea is to view the value stored in a variable not as a function of the current state, but as a function of the *relevant assignment statements*. This allows us to avoid maintaining a representation of the entire program state, instead only producing constraints for expressions which are actually required.

The particular application we consider for our new technique is the tool introduced in [7], which aims to provide Java programmers with more convenient access to optimisation technology. The tool allows an optimisation problem to be expressed in simulation form, as Java code which computes the objective value given the decisions. This code could be used directly to solve the optimisation problem by searching over possible combinations of decisions and comparing the computed results, but this is likely to be very inefficient. Instead, the tool in [7]

translates the simulation code into constraints, and then uses an *off-the-shelf* CP solver to find a set of decisions resulting in the optimal return value.

Experimental results using examples from this tool demonstrate that our new technique for modelling destructive assignments is superior to previous approaches, and can produce optimisation models comparable in efficiency to a simple hand written model for the same problem.

2 Running Example

As a running example throughout the paper we consider a smartphone app for group pizza ordering. Each member of the group nominates a number of slices and some ingredient preferences. The app automatically generates a joint order of minimum cost which provides sufficient pizza for the group, assuming that a person will only eat pizza with at least one ingredient they like and no ingredients they dislike. After approval from the user, the order is placed electronically.

Our focus is on the optimisation aspect of the application: finding the cheapest acceptable order. We assume that for each type of pizza both a price per pizza and a price per slice is specified. The order may include surplus pizza if it is cheaper to buy a whole pizza than the required number of individual slices.

Figure 1 shows a Java method defining this optimisation problem, called `buildOrder`. The problem parameters are the contents of the `people` list and the details stored in the `menu` object when `buildOrder` is called. Each call to the method `choiceMaker.chooseFrom` indicates a decision to be made, where the possible options are the objects included in the list `pizzas`. The objective is to minimise the return value, which is the total cost of the order.

3 Translating Code into Constraints

To evaluate different possible translations from procedural code to constraints we use examples from the tool in [7]. This tool actually performs the translation on demand at run-time (not as a compile time operation), which complicates the translation process somewhat. For the purpose of this paper however we will ignore such implementation details, using the following abstraction to simplify the description of the different translations.

We consider the translation to be split into two phases. In the first phase the code is flattened into a linear sequence of assignment statements, each of which has some conditions attached. We describe this transformation briefly below. In the second phase, which is the main focus of the paper, the flattened sequence of assignments is translated into constraints.

3.1 Flattening

In the Java programming language only the assignment statement changes the state of the program. All other constructs simply influence which other statements will be executed. It is therefore possible to emulate the effect of a piece of Java code using a sequence of assignment statements, each with an attached set

```

int buildOrder() {
  order = new Order(menu);
  for(Person person : people) {
    // Narrow down acceptable pizzas
    pizzas.clear();
    for(OrderItem item : order.items)
      if(person.willEat(item))
        pizzas.add(item);
    // Choose from these for each slice
    for(int i = 0; i < person.slices; i++) {
      OrderItem pizza =
        choiceMaker.chooseFrom(pizzas);
      pizza.addSlice();
    } }
  return order.totalCost();
}

class Order {
  List<OrderItem> items;
  int totalCost() {
    int totalcost = 0;
    for(OrderItem item : items)
      totalcost += item.getCost();
    return totalcost;
  }
}

class OrderItem
{
  int pricePerPizza;
  int pricePerSlice;
  int fullPizzas = 0;
  int numSlices = 0;

  void addSlice() {
    numSlices = numSlices + 1;
    if(numSlices == slicesPerPizza) {
      numSlices = 0;
      fullPizzas = fullPizzas + 1;
    } }

  int getCost() {
    int cost = fullPizzas * pizzaPrice;
    if(numSlices > 0) {
      int slicesCost =
        numSlices * slicePrice;
      if(slicesCost > pizzaPrice)
        slicesCost = pizzaPrice;
      cost = cost + slicesCost;
    } }
}

```

Fig. 1: A Java simulation of a pizza ordering optimisation problem.

of conditions controlling whether or not it should be executed. The conditions reflect the circumstances under which this statement would be reached during execution of the original code.

The flattening process involves unrolling loops³, substituting method bodies for method calls, and removing control flow statements after adding appropriate execution conditions for the child statements. As an example, consider the method `getCost` shown in Figure 1. To flatten an `if` statement we simply add the `if` condition to the execution conditions of every statement within the `then` part. The body of `getCost` can be flattened into the sequence of conditional assignments given below.

Conditions	Variable	Assigned Value
1.		<code>cost := fullPizzas × pizzaPrice</code>
2. <code>(numSlices > 0)</code>	<code>: slicesCost</code>	<code>:= numSlices × slicePrice</code>
3. <code>(numSlices > 0, slicesCost > pizzaPrice)</code>	<code>: slicesCost</code>	<code>:= pizzaPrice</code>
4. <code>(numSlices > 0)</code>	<code>: cost</code>	<code>:= cost + slicesCost</code>

³ The tool in [7] only supports loops with exit conditions unaffected by the decisions, or iteration over bounded collections. This means the number of loop iterations is always bounded. For unbounded loops partial unrolling can be performed, and the final model will be an under-approximation of the behaviour of the program.

An important optimisation is to consider the declaration scope of the variable being assigned to. For example, if a variable is declared inside the `then` part of an if statement (as is the case for the `slicesCost` variable above), assignments to that variable need not depend on the if condition. In any execution of the original code where the if condition does not hold, variables declared within the `then` part would not be created, and therefore their values are irrelevant. This means assignments 2 and 3 above do not need the condition `numSlices > 0`.

So far we have only considered assignments to local variables. For assignments to object fields we record the variable holding the object as well as the field identifier. Note that the value of the object variable may depend on the decisions, so the concrete object whose field is updated by a particular assignment statement is not necessarily known.

Some branch conditions may be independent of the decisions. In this case the branch which doesn't apply can be deleted and the condition removed. Similarly, expressions which do not depend on the decisions can be calculated upfront during the flattening process. In our running example, the code used to find acceptable order items for each person does not depend on any decisions, so rather than including assignments originating in this part of the code in the flattened sequence, we simply calculate these lists and then use them as constants.

Figure 3 in the appendix shows the sequence of assignments produced by flattening our example function `buildOrder`. Note that calls to `ChoiceMaker` methods are left untouched—these represent the creation of new decision variables.

4 Modelling Assignments: Existing Techniques

Using the flattening transformation described above and a straight forward translation of mathematical and logical expressions, we reduce the problem of representing Java code by constraints to that of modelling (conditional) assignment statements. In this section we describe two existing approaches to this, while in the next section we introduce a new proposed approach.

4.1 Typical CP Approach

One obvious technique for modelling assignments, and that used in [7, 2, 4], is to create a new version of the assigned-to variable for each assignment, and then use the latest version whenever a variable is referred to as part of an expression. If the assignment has some conditions, the new version of the variable can be constrained to equal either the assigned value or the previous version, depending on whether or not the conditions hold. This is easily achieved using a pair of implications, or alternatively using an `element` constraint with the condition as the index. The `element` constraint has the advantage that some propagation is possible before the condition is fixed, so we will use this translation.

The constraint arising from a local variable assignment is shown below, where `localvar0` is the latest version of `localvar` before the assignment, and `localvar1` is the new variable that results from the assignment, which will become the new latest version of `localvar`. For convenience, in the rest of the paper we use a simplified syntax for these constraints (also shown below).

assignment: $\text{condition} : \text{localvar} := \text{expression}$
 \downarrow
constraint: $\text{element}(\text{bool2int}(\text{condition})+1, [\text{localvar0}, \text{expression}], \text{localvar1})$
simple syntax: $\text{localvar1} = [\text{localvar0}, \text{expression}][\text{condition}]$

This translation is only correct for local variables. Field assignments are more difficult to handle due to the possibility of aliasing between objects. However, if the set of concrete objects which may be referred to by an object variable is finite (which is the case for our application), then it is possible to convert all field assignments into equivalent assignments over local variables, after which the translation above can be applied.

For each concrete object, a local variable is created to hold the value of each of its fields. In the following we name these variables using the object name and the field name separated by an underscore. Then every field assignment is replaced by a sequence of local variable assignments, one for each of the possibly affected concrete objects. These new assignments retain the original conditions, and each also has one further condition: that its corresponding concrete object is the one referred to by the object variable. Where necessary to avoid duplication, an intermediate variable is created to hold the assigned expression.

An example of this conversion is shown below, where `Obj1`, `Obj2`, and `Obj3` are the concrete objects possibly referred to by `objectvar` at this time.

field assignment: $\text{condition} : \text{objectvar}.\text{field} := \text{expression}$
 \downarrow
assignments: $\text{condition} \wedge (\text{objectvar} = \text{Obj1}) : \text{Obj1}.\text{field} := \text{expression}$
 $\text{condition} \wedge (\text{objectvar} = \text{Obj2}) : \text{Obj2}.\text{field} := \text{expression}$
 $\text{condition} \wedge (\text{objectvar} = \text{Obj3}) : \text{Obj3}.\text{field} := \text{expression}$

The final requirement is to handle references to object fields. Recall that we know the set of concrete objects which may be referred to by the object variable, and we have a local variable storing the field value for each of these concrete objects. Therefore field references can be represented by a pair of element constraints sharing an index as shown below, where `fieldrefvar` is an intermediate variable representing the retrieved value. We assume again that `objectvar` may refer to `Obj1`, `Obj2` or `Obj3`.

field reference: $\text{objectvar}.\text{field}$
 \downarrow
constraints: $\text{element}(\text{indexvar}, [\text{Obj1}, \text{Obj2}, \text{Obj3}], \text{objectvar})$
 $\text{element}(\text{indexvar}, [\text{Obj1}.\text{field}, \text{Obj2}.\text{field}, \text{Obj3}.\text{field}], \text{fieldrefvar})$

In summary, this approach involves two steps. First the list of assignments is modified to replace field assignments with equivalent local variable assignments, introducing new variables as required. Then the new list (now containing only local variable assignments) is translated into constraints, with special handling for field references. This approach is quite simple, but can result in a very large model if fields are used extensively. To see the result of applying this translation to a portion of our running example, see Figure 2(b).

4.2 Typical SMT Approach

One of the main reasons for the significant advances in program symbolic reasoning (e.g. verification and testing) during the last decade has been the remarkable progress in modern SMT solvers (we refer the reader to [1] for details). When using SMT, local variable assignments can be translated in the same way as for the CP approach (adding a new version of the variable for each assignment), but using an if-then-else construct (*ite* below) instead of an element constraint.

assignment: condition : localvar := expression
 ↓
formula: localvar1 = ite(condition, expression, localvar0)

For field assignments, it is more convenient to use the theory of arrays. This theory extends the theory of uninterpreted functions with two interpreted functions *read* and *write*. McCarthy proposed [10] the main axiom for arrays:

$$\begin{aligned}
&\forall a, i, j, x \text{ (where } a \text{ is an array, } i \text{ and } j \text{ are indices and } x \text{ is a value)} \\
& i = j \rightarrow \text{read}(\text{write}(a, i, x), j) = x \\
& i \neq j \rightarrow \text{read}(\text{write}(a, i, x), j) = \text{read}(a, j)
\end{aligned}$$

Note that since we are not interested in equalities between arrays we only focus on the non-extensional fragment.

Following the key idea of Burstall [3] and using the theory of arrays, we define one array variable for each object field. An assignment to a field is modelled as a write to an array containing the field values indexed by an object. The result is a new array variable representing the new state of the field for all relevant objects. This is much more concise and efficient than creating an explicit new variable for each concrete object. Moreover, this modelling syntactically encodes the fact that two fields cannot be aliased. Therefore, whenever one field is updated, only the corresponding array variable is modified and any other field is unchanged.

We still need to handle assignments with conditions. If the condition does not hold all field values should remain the same, so we can simply use an *ite* to ensure that in this case the new array variable is equal to the previous version.

field assignment: cond : objectvar.field := expression
 ↓
formula: field1 = ite(cond, *write*(field0, objectvar, expression), field0)

A reference to an object field is represented as a read of the latest version of the field array, using the object variable as the lookup index.

field reference: objectvar.field
 ↓
formula: *read*(field0, objectvar)

For a more complete example, see Figure 2(c). This example clearly demonstrates that the SMT formula can be much more concise than the CP model arising from the translation discussed in the previous section. Its weakness is its inability to reason over disjunction (compared to *element* in the CP approach). The approaches are compared further in Section 5.2.

<pre> pizzas1 := order items acceptable to 1st person pizza1 := chooseFrom(pizzas1) pizza1 . numSlices := pizza1.numSlices + 1 b1 := pizza1.numSlices == slicesPerPizza (b1) : pizza1 . numSlices := 0 </pre> <p style="text-align: center;">(a) First Five Assignments</p>
<pre> pizza1 in {Veg, Marg} element(index1, [Veg,Marg], pizza1) element(index1, [Veg_numSlices0,Marg_numSlices0], pizza1_numSlices0) temp1 = pizza1_numSlices0 + 1 Marg_numSlices1 = [Marg_numSlices0,temp1][pizza1 = Marg] Veg_numSlices1 = [Veg_numSlices0,temp1][pizza1 = Veg] element(index2, [Veg,Marg], pizza1) element(index2, [Veg_numSlices1,Marg_numSlices1], pizza1_numSlices1) b1 = (pizza1_numSlices1 == slicesPerPizza) Marg_numSlices2 = [Marg_numSlices1,0][b1 ^ pizza1 = Marg] Veg_numSlices2 = [Veg_numSlices1,0][b1 ^ pizza1 = Veg] </pre> <p style="text-align: center;">(b) CP Translation: Constraints</p>
<pre> (pizza1 = Marg) ∨ (pizza1 = Veg) numSlicesArray1 = write(numSlicesArray0, pizza1, read(numSlicesArray0,pizza1)+1) b1 = (read(numSlicesArray1,pizza1) = slicesPerPizza) numSlicesArray2 = ite(b1, write(numSlicesArray1,pizza1,0), numSlicesArray1) </pre> <p style="text-align: center;">(c) SMT Translation: Formula</p>

Fig. 2: (a) The first 5 assignments in the flattened sequence for the running example, and the constraints arising from applying (b) the obvious CP approach, or (c) the SMT approach to translation. We assume an instance where the acceptable pizzas for the first person are Veg(etarian) and Marg(harita).

5 A New Approach to Modelling Assignments

The main problem with the CP approach presented earlier is the excessive number of variables created to store new field values for every object possibly affected by a field assignment. Essentially we maintain a representation of the complete state of the program after each execution step.

This is not actually necessary. Our only real requirement is to ensure that the values retrieved using field references are correctly determined by the assignment statements. Maintaining the entire state is a very inefficient way of achieving this, since we may make several assignments to a field using different object variables before ever referring to the value of that field for a particular concrete object.

We can see an example of this in our running example, where there are five assignments to the `fullPizzas` field, all of which occur before we first look up the value of this field for the concrete object `item1` (see Figure 3). There is no need to create variables to hold the value of `item1.fullPizzas` after each of the assignments. We only need to know the value after all five.

To take advantage of this observation, we move away from the state-based representation, instead simply creating a variable for each field reference, and constraining this to be consistent with the relevant assignments. The relevant assignments are those using the correct field identifier which occur before the read and may refer to the same object.

As an example, let us consider the reference to `item1.fullPizzas` mentioned above (line 30 in Figure 3). Below is the list of all assignments to any `fullPizzas` field occurring before this reference. This is a superset of the relevant assignments, which we will narrow down by reasoning about the possible values taken by the object variables.

	Conds	Object	Field/Var	Assigned Value
6.	(b1)	: pizza1	. fullPizzas	:= pizza1.fullPizzas + 1
11.	(b2)	: pizza2	. fullPizzas	:= pizza2.fullPizzas + 1
17.	(b3)	: pizza3	. fullPizzas	:= pizza3.fullPizzas + 1
22.	(b4)	: pizza4	. fullPizzas	:= pizza4.fullPizzas + 1
27.	(b5)	: pizza5	. fullPizzas	:= pizza5.fullPizzas + 1

Each of the `pizza` variables is a slice choice which must correspond to one of the acceptable order items for the relevant person. We have assumed previously that the acceptable pizzas for person 1 are `Veg` and `Marg`. Let us assume that person 2 will accept either `Marg` or `Mush(room)`, and also that the order items are listed in the order `[item1,item2,item3] = [Veg,Marg,Mush]`. Given these assumptions, the possible values for each relevant object variable are shown below.

<code>pizza1</code> \in { <code>Veg,Marg</code> }	<code>pizza3</code> \in { <code>Marg,Mush</code> }	<code>pizza5</code> \in { <code>Marg,Mush</code> }
<code>pizza2</code> \in { <code>Veg,Marg</code> }	<code>pizza4</code> \in { <code>Marg,Mush</code> }	<code>item1</code> = <code>Veg</code>

Since `pizza3`, `pizza4` and `pizza5` cannot possibly refer to the same object as `item1`, the last three assignments in the table above cannot affect the value of `item1.fullPizzas` and so are not relevant. If the list included an assignment which definitely applied (an unconditional assignment with object known to equal `item1`) then we could also remove any assignments earlier than that. When there is no such assignment (as in our case), we need to allow for the possibility that none of the assignments apply to `item1`, in which case the value of `item1.fullPizzas` will be the initial value stored in the field. To handle this we add an extra unconditional assignment at the beginning of the list to set `item1.fullPizzas` to its original value (in our case this original value is 0). Note that if we had already created constraints for an earlier read of `item1.fullPizzas`, we should avoid creating duplicate expressions and constraints by using the result of that read as the initial value and only considering assignments after that point. Our final list of relevant assignments is as follows.

<code>item1</code> . <code>fullPizzas</code> := 0
(b1) : <code>pizza1</code> . <code>fullPizzas</code> := <code>pizza1.fullPizzas</code> + 1
(b2) : <code>pizza2</code> . <code>fullPizzas</code> := <code>pizza2.fullPizzas</code> + 1

To make the model correct, we need constraints ensuring that the retrieved value (`item1.fullPizzas`) corresponds to the value associated with the most recent of these assignments which applied to the correct object and actually took place

(had a true execution condition). To achieve this we introduce a new integer variable `indexvar` whose value indicates which of the relevant assignments is the one corresponding to this read. We use three element constraints to ensure that the assignment at the selected index applies to the read object and is actually executed, and that the retrieved value is equal to the value assigned at this index.

```

element(indexvar, [item1,pizza1,pizza2], item1)
element(indexvar, [true,b1,b2], true)
element(indexvar, [0,pizza1_fullPizzas + 1,pizza2_fullPizzas + 1], item1_fullPizzas)

```

Note that `pizza1_fullPizzas` and `pizza2_fullPizzas` are the variables introduced for the field references used as part of the assigned values. These would be constrained using their own list of relevant assignments.

The only remaining requirement is that we must choose the *latest* applicable assignment. Using the natural order for the arrays this corresponds to the greatest index. We therefore add additional constraints stating that if at index i the object variables are equal and the execution condition is true, then the selected index must be greater than $i - 1$.

```

(b1 ∧ pizza1 = item1) → indexvar > 1
(b2 ∧ pizza2 = item1) → indexvar > 2

```

The general form of the constraints used for field references is shown below. References to local variables are treated as field references where the object variable used is the same as that used for all relevant assignments. When this is the case the first element constraint is not required (as it is trivially satisfied), and the implications can be simplified to use only the execution condition on the left hand side. Other obvious simplifications are also applied.

<i>field reference:</i>	queryobj.field
	↓
<i>relevant assignments:</i>	queryobj.field := init cond1 : obj1.field := expr1 ... condn : objn.field := exprn
	↓
<i>constraints:</i>	element(indexvar, [queryobj, obj1, ..., objn], queryobj) element(indexvar, [true, cond1, ..., condn], true) element(indexvar, [init, expr1, ..., exprn], queryobj.field) (cond1 ∧ queryobj = obj1) → indexvar > 1 ... (condn ∧ queryobj = objn) → indexvar > n

In some cases, we can detect a pattern to the relevant assignments which allows us to use a more specialised constraint. The next section describes three commonly occurring special cases and their alternative treatment.

5.1 Special Cases

The three special cases we look for are Boolean variables, sequences of assignments representing a sum calculation, and sequences of assignments representing

a maximum or minimum calculation. The translation of assignments automatically detects the cases described below and uses the more efficient translation.

Boolean Variables When the referenced variable is of type `bool`, we can define a Boolean expression for the retrieved value rather than using `element` constraints and implications as shown above. Usually an `element` constraint has an advantage over a disjunction of equalities because it may reduce the domain of the value when there are still several possibly applicable array elements with different domains. However, if the value is a Boolean, then no propagation is possible until all possibly applicable array elements have the same value.

For a reference to a `bool` field we constrain the result to be equal to the expression shown below, which is true if some assignment with a true value applies and no later assignment with a false value applies.

$$\begin{array}{l}
 \textit{field reference: } q.\textit{field} \\
 \textit{assignments: } c_i : o_i.\textit{field} := e_i \quad i \in 1..n \\
 \textit{expression: } \bigvee_{i \in 1..n} \left(c_i \wedge e_i \wedge (o_i = q) \wedge \bigwedge_{j \in i+1..n} (e_j \vee \neg c_j \vee (o_j \neq q)) \right)
 \end{array}$$

Local variables are handled in the same way except without the object equalities. Using this constraint instead of the generic constraint eliminates the need to introduce an index variable, and allows simplifications to be performed when objects are known to be equal, assigned values are fixed, or assignments are unconditional.

Sum Calculations Computations often involve taking the sum of a set of numbers. In a procedural language, sums are commonly calculated by iteratively adding each number to a variable representing the total. This coding pattern results in a sequence of writes where each written value is an addition of the previous value of this variable and some other number. When this pattern is detected, we can replace the usual constraints with a `sum` constraint.

$$\begin{array}{l}
 \textit{relevant assignments: } \textit{total} := 0 \\
 \quad \textit{total} := \textit{total} + \textit{value1} \\
 \quad \textit{total} := \textit{total} + \textit{value2} \\
 \quad \textit{total} := \textit{total} + \textit{value3} \\
 \quad \quad \downarrow \\
 \textit{constraint: } \quad \textit{finaltotal} = \textit{sum}([\textit{value1}, \textit{value2}, \textit{value3}])
 \end{array}$$

In the example above all assignments were unconditional and used exactly the same variable (the local variable `total`). It is also possible to use a general form of this constraint for sequences of assignments which do not represent a pure sum but a related calculation, such as counting the objects which satisfy some condition. Consider again the relevant writes for the reference to `item1.fullPizzas` discussed earlier. A better constraint for `item1.fullPizzas` is shown below.

```

assignments:      item1 . fullPizzas := 0
                    (b1) : pizza1 . fullPizzas := pizza1.fullPizzas + 1
                    (b2) : pizza2 . fullPizzas := pizza2.fullPizzas + 1
                                ↓
constraint:      item1_fullPizzas =
                    sum([bool2int(b1 ∧ pizza1=item1), bool2int(b2 ∧ pizza2=item1)])

```

The biggest advantage of this constraint is that it removes the need to create and constrain the variables `pizza1_fullPizzas` and `pizza2_fullPizzas`. These can be excluded from the model entirely as they were only used to build a new expression to re-assign to the same variable, and are now not required to define the values retrieved from this field.

The general form of the alternative constraint used for sums is given below. This constraint is applicable whenever for all relevant assignments, either the assignment definitely applies (this will be the first assignment in the list), or the assigned value is an addition with one argument being a reference to the assigned-to variable.

```

field reference: queryobj.field
assignments:    condi : obji.field := obji.field + expri    i ∈ 1..n
constraint:      q.field = sum([expri × bool2int(condi ∧ obji = queryobj) | i ∈ 1..n])

```

Max/Min Calculations Another common coding pattern is to calculate a maximum or minimum by iterating through a list of values overwriting a variable each time a smaller/larger value is found. As with sum, we can detect this pattern when building the constraints for the final read of the variable. This time in order for the alternative constraint to apply, every non-initialisation assignment must have a condition which compares the current value of the variable with the assigned value. When there are no other assignment conditions, and the variable is a local variable (or the assigned-to object is known to equal the read object for all relevant assignments), we can use a max/min constraint as shown below.

```

assignments:      max := init
                    (value1 > max) : max := value1
                    (value2 > max) : max := value2
                                ↓
constraint:      finalmax = max([init, value1, value2])

```

We can again extend to the general case for field assignments where the object variables used for the writes and read are not known to be equal, and handle assignments with additional conditions. When extra conditions are present, we are calculating the maximum or minimum value for which these additional conditions hold. For a maximum, we constrain the result to be no less than any value for which the extra conditions hold, and to equal one of the values for which the conditions hold. Minimum is handled equivalently.

```

field reference: queryobj.field
assignments:    (condi ∧ valuei > obji.field) : obji.field := valuei    i ∈ 1..n
                                ↓
constraints:      ∨i∈1..n condi ∧ (obji = queryobj) ∧ (result = valuei)
                    ∧i∈1..n (condi ∧ obji = queryobj) → result ≥ valuei

```

Table 1: Comparing three approaches to modelling destructive assignments.

Benchmark	smt	Time (secs)						Failures (000s)					
		orig	orig+sp	new	new+sp	hand	orig	orig+sp	new	new+sp	hand		
proj1	200	2.2	23.0	0.1	12.1	0.1	0.1	56	0	34	0	0	
	225	2.4	3.2	0.1	1.5	0.1	0.1	9	0	4	0	0	
	250	1.6	61.9 3	0.1	61.7 3	0.1	0.1	99	0	127	0	0	
proj2	22	115.6 2	84.8 1	42.7 1	51.7 2	23.7 1	7.6	39	31	110	35	22	
	24	221.1 7	286.9 9	167.6 5	170.9 6	129.2 4	92.2 4	92	89	368	239	280	
	26	262.7 8	376.2 16	293.3 10	255.9 11	137.9 6	128.9 5	120	144	583	251	452	
pizza	3	56.0	37.4 1	25.1	7.0	3.1	2.0	175	118	30	14	0	
	4	226.4 8	180.9 7	175.7 7	138.0 4	79.3 2	2.1	544	541	377	252	1	
	5	480.9 22	411.8 18	407.5 18	343.4 13	298.3 12	2.2	1170	1216	865	945	7	

5.2 Comparison with Earlier Approaches

We compared the three presented translation techniques experimentally, using the pizza ordering example plus two benchmarks used in [7] (those that don't require collections). We used 30 instances for each of several different sizes to evaluate scaling behaviour. For the **original** and **new** CP approaches we show the effect of adding **+special** cases. Special cases can be detected in the original method, but only for local variables. Using the new translation makes these cases also recognisable for fields. As a reference we also include a fairly naive **hand** written model for each problem. The Java code defining the problems and all compared constraint models are available online at www.cs.mu.oz.au/~pjs/optmodel.

The CP models were solved using the lazy clause generation solver Chuffed. The SMT versions were solving using Z3 [5]. Z3, like most SMT solvers, does not have built-in support for optimisation. We used a technique similar to [11] to perform optimisation using SMT: in incremental mode, we repeatedly ask for a solution with a better objective value until a not satisfiable result is returned.

Table 1 shows average time to solve and failures for the different models. The small number next to the time indicates number of timeouts (> 600s). The results show while the SMT approach does compete with the original approach, with special case treatment it does not. The **new** approach is quite superior and in fact has a synergy with special cases (since more of them are visible). **new+sp** competes with **hand** except for **pizza** where it appears that the treatment of the relationship between slices and full pizzas used in **hand** is massively more efficient than the iterative approach in the simulation.

6 Collection Operations

The code for the pizza ordering example makes use of collection classes from the Java Standard Library: **Set**, **List** and **Map**. In this case no special handling is required as all collection operations are independent of the decisions, but often it is more natural to write code where that is not the case. For example, say we wished to extend our application to choose between several possible pizza

outlets, each with a different menu. We could do this by adding one extra line at the beginning of the `buildOrder` function.

```
menu = chooseFrom(availableMenus);
```

This change means the contents of the `OrderItem` list in the `Order` class will depend on the decisions, so the `for` loop iterating over this list (in `buildOrder`) will perform an unknown (though bounded) number of iterations, and the result of any query operations on this list will also depend on the decisions.

In [7], collection operations were supported by introducing appropriately constrained variables representing the state of each collection after each update operation (e.g. `List.add`). Query operations (e.g. `List.get`) were represented as a function of the current state of the relevant collections. This is analogous to the way field assignments were handled, with the same drawbacks.

Fortunately our new technique can also be extended to apply to collection operations, resulting in a much more efficient representation. Where previously the flattened list of state changing operations contained only assignments, we now also include collection update operations. Then every query operation on a collection is treated analogously to a field reference. That is, a new variable is created to hold the returned value, and constraints are added to ensure that this value is consistent with the relevant update operations.

Below we provide details of the constraints used for `List` operations. `Set` and `Map` operations are treated similarly; a detailed description is omitted for brevity. We then give experimental results using collection-related benchmarks from [7].

6.1 Example: List

For the `List` class we support the following operations.

update operations: `add` (at end of list), `replace` (item at index)
query operations: `get` (item at index), `size`

We will explain how these operations are supported by example. A code snippet containing one of each operation type is shown below. Also shown are the assumed possible variable values and initial list contents, and the flattened list of collection update operations. Each update operation has an associated condition, list, index and item. For the `add` operation, the index is a variable `size1` holding the current size of `list1`. The first three operations in the table reflect the original contents of the lists.

<pre>if(cond) { list1.add(A); list1.replace(0, item); }</pre>	<pre>list1 ∈ {L1,L2,L3} ind ∈ {0,1,2} list2 ∈ {L1,L2,L3} item ∈ {A,B,C}</pre>	<table border="1"> <thead> <tr> <th>Cond</th> <th>List</th> <th>Index</th> <th>Item</th> </tr> </thead> <tbody> <tr> <td></td> <td>L1</td> <td>[0]</td> <td>:= A (<i>add</i>)</td> </tr> <tr> <td></td> <td>L1</td> <td>[1]</td> <td>:= B (<i>add</i>)</td> </tr> <tr> <td></td> <td>L2</td> <td>[0]</td> <td>:= C (<i>add</i>)</td> </tr> <tr> <td>cond : list2</td> <td>list1</td> <td>[size1]</td> <td>:= A (<i>add</i>)</td> </tr> <tr> <td>cond : list1</td> <td>list1</td> <td>[0]</td> <td>:= item (<i>repl</i>)</td> </tr> </tbody> </table>	Cond	List	Index	Item		L1	[0]	:= A (<i>add</i>)		L1	[1]	:= B (<i>add</i>)		L2	[0]	:= C (<i>add</i>)	cond : list2	list1	[size1]	:= A (<i>add</i>)	cond : list1	list1	[0]	:= item (<i>repl</i>)
Cond	List	Index	Item																							
	L1	[0]	:= A (<i>add</i>)																							
	L1	[1]	:= B (<i>add</i>)																							
	L2	[0]	:= C (<i>add</i>)																							
cond : list2	list1	[size1]	:= A (<i>add</i>)																							
cond : list1	list1	[0]	:= item (<i>repl</i>)																							
(a) Code	(b) Variables	(c) Update Operations																								

With our limited set of supported update operations (which is nevertheless sufficient to cover all code used in the benchmarks from [7]), the size of a list is simply the number of preceding `add` operations applying to this list and having true execution conditions. Note that the `replace` operation is not relevant to `size`.

```
query:      sizeresult := list2.size()
constraint: sizeresult = sum([ bool2int(list2=L1), bool2int(list2=L1),
                               bool2int(list2=L2), bool2int(list2=list1^cond) ])
```

A `get` query is treated almost exactly like a field reference. The value returned must correspond to the most recent update operation with true execution condition which applied to the correct list and index. There is however one extra complication to be considered. A `get` operation which is not actually executed should not impose any constraint on its index and list variables. In the constraints shown below, the final element of each array has been added to allow the result variable for our `get` query to take the arbitrary value `A` when the `get` would not be executed (`size2>ind` is false). Without this the constraints would incorrectly force `ind` and `list2` to correspond to an executed update operation.

```
query:      getresult := list2.get(ind)
constraints: element(indexvar, [L1,L1,L2,list1,list1,list2], list2)
               element(indexvar, [0,1,0,size1,0,ind], ind)
               element(indexvar, [true,true,true,cond,cond,¬(sizeresult>ind)], true)
               element(indexvar, [A,B,C,A,item,A], getresult)
               (list2=L1) ^ (ind=1)           → indexvar ≥ 2
               (list2=list1) ^ (ind=size1) ^ cond → indexvar ≥ 4
               (list2=list1) ^ (ind=0) ^ cond  → indexvar ≥ 5
               ¬(sizeresult>ind)              → indexvar ≥ 6
```

6.2 Comparison on Benchmarks with Collections

Table 2 compares the original and new translation approaches with hand written models, using problems involving variable collections from [7]. It is clear that the new translation substantially improves on the old in most cases, and is never very much worse (`bins,knap1`). With the addition of special case treatment the new translation is often comparable to the hand written model, though certainly not always (`proj3, route`). In a few instances it is superior (`bins,golf`), this may be because it uses a sequential search extracted from the Java problem structure, or indeed that the intermediate variables it generates give more scope for reusable nogood learning.

7 Conclusion

Effective modelling of destructive assignment is essential for any form of reasoning about procedural code. We have developed a new encoding of assignment and state that gives effective propagation of state-related information. We demonstrate the effectiveness of this encoding for the automatic generation of optimisation models from simulation code, showing that the resulting model can be comparable in efficiency to a hand-written optimization model.

Table 2: Comparison on examples which use variable collections.

Benchmark	Time (secs)					Failures (000s)			
	orig+sp	new	new+sp	hand		orig+sp	new	new+sp	hand
bins	12	2.6	5.3	1.1	1.2	8.1	32.5	6.2	13.8
	14	82.8 ¹	129.6 ³	7.6	18.0	95.4	612.9	75.1	169.9
	16	327.2 ¹⁵	391.6 ¹⁵	84.8	141.6 ⁵	315.1	1617.6	749.5	1355.0
golf	4,3	0.7	0.2	0.2	21.3	0.7	0.8	0.7	159.7
	4,4	3.4	2.0	0.3	0.1	0.8	6.7	0.0	0.0
	5,2	2.4	0.8	0.3	1.5	0.4	0.3	0.0	12.3
golomb	8	1.3	1.2	1.2	1.2	10.8	10.4	10.4	24.0
	9	14.0	12.9	12.9	13.7	55.4	51.9	51.9	149.4
	10	161.5	144.1	151.8	178.8	281.1	284.5	284.5	1211.0
knap1	70	2.1	8.3	2.8	1.8	33.3	2.2	1.9	33.3
	80	7.5	18.4	7.1	6.8	95.7	3.5	3.5	95.7
	90	14.2	31.9	12.7	13.9	180.2	4.5	4.5	180.2
knap2	70	20.9	23.2	22.4	34.7	247.8	245.4	245.4	425.8
	80	88.4 ²	87.7 ²	93.9 ²	117.5 ³	935.2	901.8	915.0	1253.1
	90	223.6 ⁵	229.9 ⁵	230.9 ⁵	207.0 ⁵	2263.9	2182.7	2199.3	2085.5
knap3	40	26.2	0.9	0.3	0.2	14.3	0.5	0.4	1.3
	50	81.1	2.2	1.3	0.1	25.0	0.8	0.6	2.4
	60	295.2 ⁶	4.2	1.8	0.4	58.7	1.4	1.2	10.2
proj3	10	153.9 ⁵	2.3	2.4	0.1	289.3	9.3	11.6	0.1
	12	509.4 ²⁴	28.0	20.7	0.1	778.5	83.5	92.1	0.2
	14	600.0 ³⁰	133.9 ²	102.9 ¹	0.1	807.3	299.5	394.5	0.5
route	5	34.2	1.7	1.7	0.2	34.0	6.3	6.3	2.3
	6	338.3 ³	43.7	43.1	0.8	195.8	57.5	57.5	7.6
	7	600.0 ³⁰	536.9 ²⁰	502.5 ¹⁷	2.7	263.2	286.9	333.1	19.2
talent	3,8	11.1	3.4	0.9	0.8	25.9	17.2	5.0	8.9
	4,9	170.8	42.7	8.8	7.3	159.7	127.3	31.1	52.4
	4,10	545.5 ²²	223.0 ¹	77.9	54.6	459.7	510.8	178.1	212.5

In the future we will investigate the use of this encoding for symbolic execution in applications such as test generation. The main difference for these applications is the lack of a known initial state. This will require the creation of variables to represent unknown initial field values, with constraints ensuring that if a pair of object variables are equal then their corresponding initial field variables are also equal.

Uncertainty about the initial state will also have an impact on the number of relevant assignments for each field reference. For a query object with unbounded domain all assignments to the same field occurring prior to the read are relevant, unless one of these is an unconditional assignment using this exact variable.

These differences may mean that redundant constraints relating reads to each other (which we have not discussed in this paper due to their lack of impact for our application) become more important for effective propagation.

Acknowledgments NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council.

References

1. Armin Biere, Marijn J. H. Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability*, volume 185. February 2009.
2. A. Brodsky and H. Nash. CoJava: optimization modeling by nondeterministic simulation, in constraint programming. In *Principles and Practice of Constraint Programming (CP)*, pages 91–107, 2006.
3. Rod Burstall. Some techniques for proving correctness of programs which alter data structures. *Machine Intelligence*, 7:23–50, 1972.
4. H. Collavizza, M. Rueher, and P. Van Hentenryck. CPBPV: a constraint-programming framework for bounded program verification. *Constraints*, 15(2):238–264, 2010.
5. Leonardo Mendonça de Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *TACAS*, pages 337–340, 2008.
6. R. W. Floyd. Assigning meanings to programs. In *Proceedings of the American Mathematical Society Symposia on Applied Mathematics*, volume 19, pages 19–31, 1967.
7. K. Francis, S. Brand, and P.J. Stuckey. Optimization modelling for software developers. In M. Milano, editor, *Proceedings of the 18th International Conference on Principles and Practice of Constraint Programming*, number 7514 in LNCS, pages 274–289. Springer, 2012.
8. C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, October 1969.
9. James C. King. Symbolic Execution and Program Testing. *Com. ACM*, pages 385–394, 1976.
10. John McCarthy. Towards a mathematical science of computation. In *IFIP Congress*, pages 21–28, 1962.
11. Roberto Sebastiani and Silvia Tomasi. Optimization in SMT with $LA(Q)$ cost functions. In *IJCAR*, pages 484–498, 2012.

A Assignments for Running Example

	Conds	Object	Field/Var	Assigned Value
1.			pizzas1	:= <i>order items acceptable to 1st person</i>
2.			pizza1	:= chooseFrom(pizzas1)
3.		pizza1 .	numSlices	:= pizza1.numSlices + 1
4.			b1	:= pizza1.numSlices == slicesPerPizza
5.	(b1) :	pizza1 .	numSlices	:= 0
6.	(b1) :	pizza1 .	fullPizzas	:= pizza1.fullPizzas + 1
7-11.	<i>repeat assignments 2-6 for 2nd slice (using vars pizza2 and b2)</i>			
12.			pizzas2	:= <i>order items acceptable to 2nd person</i>
13.			pizza3	:= chooseFrom(pizzas2)
14.		pizza3 .	numSlices	:= pizza3.numSlices + 1
15.			b3	:= pizza3.numSlices == slicesPerPizza
16.	(b3) :	pizza3 .	numSlices	:= 0
17.	(b3) :	pizza3 .	fullPizzas	:= pizza3.fullPizzas + 1
18-22.	<i>repeat assignments 13-17 for 2nd slice (using vars pizza4 and b4)</i>			
23-27.	<i>repeat assignments 13-17 for 3rd slice (using vars pizza5 and b5)</i>			
28.			totalcost	:= 0
29.			item1	:= <i>first OrderItem</i>
30.			cost1	:= item1.fullPizzas × item1.pizzaPrice
31.			b6	:= item1.numSlices > 0
32.			slicesCost1	:= item1.numSlices × item1.slicePrice
33.			b7	:= slicesCost1 > item1.pizzaPrice
34.	(b7) :		slicesCost1	:= item1.pizzaPrice
35.	(b6) :		cost1	:= cost1 + slicesCost1
36.			totalcost	:= totalcost + cost1
37-44.	<i>repeat assignments 29-36 for 2nd order item (using vars item2 etc.)</i>			
45-52.	<i>repeat assignments 29-36 for 3rd order item (using vars item3 etc.)</i>			
53.			objective	:= totalcost

Fig. 3: Flattened version of buildOrder method, assuming that the list people contains two Person objects, the first requiring two slices and the second requiring three, and that the menu object lists three different types of pizza, meaning the order will contain three OrderItems. The b variables have been introduced to store branching conditions. Variables from methods called more than once and those used as the iteration variable in a loop are numbered to distinguish between the different versions.