

# Polyvariant program specialisation with property-based abstraction

John P. Gallagher

Roskilde University, Denmark

IMDEA Software Institute, Madrid, Spain

jpg@ruc.dk

In this paper we show that property-based abstraction, an established technique originating in software model checking, is a flexible method of controlling polyvariance in program specialisation in a standard online specialisation algorithm. Specialisation is a program transformation that transforms a program with respect to given constraints that restrict its behaviour. Polyvariant specialisation refers to the generation of two or more specialised versions of the same program code. The same program point can be reached more than once during a computation, with different constraints applying in each case, and polyvariant specialisation allows different specialisations to be realised. A property-based abstraction uses a finite set of properties to define a finite set of abstract versions of predicates, ensuring that only a finite number of specialised versions is generated. The particular choice of properties is critical for polyvariance; too few versions can result in insufficient specialisation, while too many can result in an increase of code size with no corresponding efficiency gains. Using examples, we show the flexibility of specialisation with property-based abstraction and discuss its application in control flow refinement, verification, termination analysis and dimension-based specialisation.

## 1 Program specialisation

Specialisation is a program transformation that transforms a program with respect to some given constraints that restrict its behaviour. A classic example is the loop in Figure 1(a) for computing  $z = x^y$ . Figure 1(b) shows the result of specialising the loop with the input constraint  $y = 3$ , unfolding the loop three times and evaluating the statement  $y--$  in the loop body.

<pre>z = 1; while (y&gt;0) {   z = x*z; y--; }</pre>	<pre>/* Input constraint y=3 */ z = 1; z = x*z; z = x*z; z = x*z;</pre>
(a)	(b)

Figure 1: Partial evaluation of a loop

Some specialisation methods could further transform the code in Figure 1(b) to  $z = x*x*x$ ; using algebraic reasoning.

As well as exploiting input constraints to specialise the program, we can perform internal specialisations based on constraints generated during program execution. For example, when specialising a statement  $\text{if}(e)\{s_1\}\{s_2\}$ , even where the test  $e$  itself cannot be evaluated, the branch  $s_1$  can be specialised with the constraint  $e$  and the branch  $s_2$  can be specialised with the constraint  $\neg e$ . This is sometimes called *driving* [24].

<pre> <b>while</b> (x&gt;0) {   <b>if</b> (y&lt;m) {     y++;   } <b>else</b> {     x--;   } } </pre> <p style="text-align: center;">(a)</p>	<pre> <b>if</b> (x&gt;0)   <b>while</b> (y&lt;m) { /* x&gt;0 */     y++; }   x--;   <b>while</b> (x&gt;0) { /* y&gt;=m */     x--; } } </pre> <p style="text-align: center;">(b)</p>
--	--

Figure 2: Polyvariant specialisation of a loop

*Polyvariant* specialisation refers to the generation of two or more specialised *versions* of the same program code. For example suppose that the statement  $\text{if}(x < 100)\{s_1\}\{s_2\}$  is reached twice during a computation, once with the constraint  $x < 100$  and the other with the constraint  $x \geq 100$ . Polyvariant specialisation gives rise to two instances of the statement in the specialised code,  $s_1$  and  $s_2$  respectively.

Figure 2 illustrates both internal specialisation and polyvariance. The “then” branch of the **if** statement in Figure 2(a) does not affect the loop condition and so if it is taken, the “then” branch is repeatedly taken until the test  $y < m$  fails. Then the “else” statement is executed ( $x--;$ ) after which the “else” branch is repeatedly taken, since it does not affect the condition  $y < m$ , until the test  $x > 0$  fails.

Thus implicitly there are two distinct loops separated by  $x--;$  and this leads to the polyvariant specialisation shown in Figure 2(b). The transformation in this case is more involved than in Figure 1; for instance, the first **while** statement does not directly correspond to the **while** statement in the original code, having a different loop test. However, the specialised program performs exactly the same sequence of computations as the original program, apart from tests that have been evaluated to true during partial evaluation and thus omitted from the specialised code. Further explanation is given in Example 3.

The main contribution of this paper is a specialisation algorithm that performs polyvariant specialisation. Instances of this algorithm have been previously used and briefly described [16, 17, 8] but these papers did not present and discuss the general algorithm. A key question is the control of polyvariance; in general there could be many (even an infinite number) of possible variants of a given program point. How does the partial evaluation algorithm determine a suitable set of variants, while ensuring termination of the specialisation algorithm.

The algorithm operates on constrained Horn clauses, which provide a representation language capable of representing the semantics of a wide range of programming languages and systems. The algorithm is parameterised by a set of properties that control the generation of polyvariance.

## 2 Preliminaries

### 2.1 Constraints and entailment

Let  $T$  be a theory and let  $\mathcal{C}^T$  be the set of formulas (also called *constraints*) constructed from the predicates and function symbols of  $T$  together with variables and boolean connectives, and the formulas true and false.  $\models_T \phi$  means that  $\phi$  is true in the theory  $T$ , where  $\phi$  is a variable-free formula. For example,  $\models_T 1 \geq 0$  where  $T$  is the theory of linear real arithmetic. In this paper, if we omit the theory subscript  $T$  we assume that  $T$  is the theory of linear real arithmetic, and we omit the symbol  $\models$  when clear from context.

Let  $\phi \in \mathcal{C}^T$  be a constraint possibly containing variables; a substitution for the variables of  $\phi$  is a *grounding substitution* if the result of applying the substitution to  $\phi$ , say  $\phi'$ , contains no variables; if  $\models_T \phi'$  the grounding substitution *satisfies*  $\phi$ .

For all constraints  $\phi$  and  $\psi$ , we say that  $\phi$  *entails*  $\psi$  in  $T$ , written  $\phi \preceq_T \psi$ , if every grounding substitution that satisfies  $\phi$  in  $T$  also satisfies  $\psi$  in  $T$ . For example,  $x \geq 1 \preceq x \geq 0$ .

We assume that there is a procedure called SAT such that for every  $\phi \in \mathcal{C}^T$ ,  $\text{SAT}(\phi)$  is true if there is some substitution that satisfies  $\phi$  and false otherwise. Using SAT, we can check entailment;  $\phi \preceq_T \psi$  if and only if  $\text{SAT}(\phi \wedge \neg\psi)$  is false.

**Definition 1 (Generalisation)** A function  $\rho : \mathcal{C}^T \rightarrow \mathcal{C}^T$  is called a generalisation operator if  $\phi \preceq_T \rho(\phi)$ .

## 2.2 Constrained Horn clause representation of programs

A constrained Horn clause (CHC) is a first-order predicate logic formula of the form  $\forall \mathbf{x}_0 \dots \mathbf{x}_k (p_1(\mathbf{x}_1) \wedge \dots \wedge p_k(\mathbf{x}_k) \wedge \phi \rightarrow p_0(\mathbf{x}_0))$ , where  $\phi$  is a finite conjunction of *constraints*,  $\mathbf{x}_0, \dots, \mathbf{x}_k$  are (possibly empty) tuples of *variables*,  $p_0, \dots, p_k$  are *predicate symbols*,  $p_0(\mathbf{x}_0)$  is the *head* of the CHC and  $p_1(\mathbf{x}_1) \wedge \dots \wedge p_k(\mathbf{x}_k) \wedge \phi$  is the *body*. Formulas of the form  $p(\mathbf{x})$  are called atomic formulas or simply *atoms*. A CHC is often written as  $p_0(\mathbf{x}_0) \leftarrow \phi, p_1(\mathbf{x}_1), \dots, p_k(\mathbf{x}_k)$  in the style of constraint logic programs, or  $p_0(X_0) \leftarrow C, p_1(X_1), \dots, p_k(X_k)$  in text form, where  $C$  is a constraint formula.

A *constrained fact* is a CHC  $p(\mathbf{x}) \leftarrow \phi$ , where  $\phi$  is a constraint over  $\mathbf{x}$ .

**Definition 2 (Ordering on sets of constrained facts)** We extend the relation  $\preceq$  to sets of constrained facts. Let  $S, S'$  be sets of constrained facts. Then  $S \preceq S'$  if for each constrained fact  $p(\mathbf{x}) \leftarrow \phi$  in  $S$  there exists a constrained fact (with variables suitably renamed)  $p(\mathbf{x}) \leftarrow \psi$  in  $S'$ , such that  $\phi \preceq \psi$ . Furthermore, if  $\rho$  is a generalisation operator on constraints then  $S \preceq \{p(\mathbf{x}) \leftarrow \rho(\phi) \mid p(\mathbf{x}) \leftarrow \phi \in S\}$ .

We do not go into detail on the translation of imperative programs to CHCs, but note that a distinct predicate symbol is generated for each program point, and the arguments of the predicate for a given program point are the values of the program variables at that point. The CHCs defining a predicate capture the transitions in an operational semantics. Figure 3 illustrates the translation of the programs in Figures 1(a) and 2(a) into CHCs. Depending on the style of semantic specification (such as small-step semantics or big-step semantics), different CHCs can be obtained for a program.

## 2.3 Constrained Horn clause derivations

The definitions of CHC derivations and partial evaluation are based on standard definitions (e.g. [20]) adapted to include constraints, and using the “resultant” style of derivation. Instead of an initial query or goal  $\leftarrow \phi, p(\mathbf{x})$ , we use with a CHC  $p(\mathbf{x}) \leftarrow \phi, p(\mathbf{x})$  and each steps resolves on a body literal.

**Definition 3 (Derivation step)** A CHC derivation step or unfolding step is defined as follows. Let  $c_1, c_2$  be CHCs, where  $c_1 = q_0(\mathbf{x}_0) \leftarrow \phi, q_1(\mathbf{x}_1), \dots, q_k(\mathbf{x}_k)$  and  $c_2 = q_i(\mathbf{y}_0) \leftarrow \phi', r_1(\mathbf{y}_1), \dots, r_m(\mathbf{y}_m)$ , with variables renamed so that  $c_1$  and  $c_2$  have no variables in common. Then the result of unfolding  $c_1$  with  $c_2$  on  $q_i(\mathbf{x}_i)$  is:

$$\left\{ \begin{array}{ll} q_0(\mathbf{x}_0) \leftarrow \phi \wedge \phi' \wedge \mathbf{x}_i = \mathbf{y}_0, \\ \quad q_1(\mathbf{x}_1), \dots, q_{i-1}(\mathbf{x}_{i-1}), \\ \quad r_1(\mathbf{y}_1), \dots, r_m(\mathbf{y}_m), \\ \quad q_{i+1}(\mathbf{x}_{i+1}), \dots, q_k(\mathbf{x}_k) & \text{if } \text{SAT}(\phi \wedge \phi' \wedge \mathbf{x}_i = \mathbf{y}_0) \\ q_0(\mathbf{x}_0) \leftarrow \text{false} & \text{otherwise} \end{array} \right.$$

<pre> start ←   p0(X, Y, Z). p0(X, Y, Z) ← Z1 = 1,   while0(X, Y, Z1). while0(X, Y, Z) ←   Y &gt; 0, Z1 = X * Z, Y1 = Y - 1,   while0(X, Y1, Z1). while0(X, Y, Z) ←   Y &lt;= 0. </pre> <p style="text-align: center;">(a)</p>	<pre> start ←   while0(X, Y, M). while0(X, Y, M) ←   X &gt; 0,   if0(X, Y, M). while0(X, Y, M) ←   X &lt;= 0. if0(X, Y, M) ←   Y &lt; M, Y1 = Y + 1,   while0(X, Y1, M). if0(X, Y, M) ←   Y &gt;= M, X1 = X - 1,   while0(X1, Y, M). </pre> <p style="text-align: center;">(b)</p>
--	--

Figure 3: CHC representation of (a) Figure 1(a) and (b) Figure 2(a)

**Definition 4 (Derivation tree)** Let  $P$  be a set of CHCs and let  $A = p(\mathbf{x}) \leftarrow \phi$  be a constrained fact over  $\mathbf{x}$ . Then a derivation tree for  $A$  in  $P$  is a tree where every node is labelled by a CHC, such that:

- the root is labelled with  $p(\mathbf{x}) \leftarrow \phi, p(\mathbf{x})$ ;
- for a non-leaf node labelled with a CHC  $c$ , its children are labelled with CHCs  $\{c_1, \dots, c_k\}$ , where  $q_i(\mathbf{x}_i)$  is some atom in the body of  $c$ ,  $\{d_1, \dots, d_k\}$  is the set of CHCs (with variables suitably renamed) in  $P$ , whose head has predicate  $q_i$ , and  $c_i$  is the result of unfolding  $c$  with  $d_i$  on  $q_i(\mathbf{x}_i)$ ;
- a leaf node is labelled by a CHC of the form  $p(\mathbf{x}) \leftarrow \phi$ , where  $\phi$  is a constraint (possibly false).

The definition is non-deterministic since any atom  $q_i(\mathbf{x}_i)$  can be selected at a non-root node.

A derivation tree can contain infinite branches. The specialisation algorithm developed in the next section depends on constructing a *partial* derivation tree, which is a finite tree following Definition 4 with the additional case that a leaf node may be labelled by a CHC of the form  $p(\mathbf{x}_0) \leftarrow \phi, q_1(\mathbf{x}_1), \dots, q_k(\mathbf{x}_k)$  ( $k > 0$ ) representing an incomplete branch of the derivation tree.

A branch of a (partial) derivation tree is a *failing* branch if it ends in a leaf labelled by  $p(\mathbf{x}) \leftarrow \text{false}$ , otherwise it is *non-failing*.

**Definition 5 (Partial evaluation of a constrained fact)** Let  $A$  be a constrained fact and  $P$  be a set of CHCs. Let  $T$  be a derivation tree for  $A$  in  $P$ . Then a partial evaluation of  $A$  in  $P$  is a finite set of CHCs  $\{c_1, \dots, c_m\}$  labelling nodes chosen from the non-root nodes of  $T$  such that there is exactly one node for each non-failing branch of  $T$ .

Clearly, the whole derivation tree does not have to be constructed in order to get a partial evaluation, but only an initial portion; then one CHC from each branch is collected.

The non-determinism in the definition of a derivation tree is resolved by an *unfolding rule*, which both selects which body atom to unfold at each step, and decides when to stop extending a branch in a (partial) derivation tree in order to return a partial evaluation.

**Definition 6 (Unfolding rule)** An unfolding rule  $U$  is a function which given a set of CHCs  $P$  and a constrained fact  $A$ , returns exactly one finite set of CHCs that is a partial evaluation of  $A$  in  $P$ . For a set  $S$  of constrained facts, the set of CHCs obtained by applying  $U$  to each element of  $S$  is called a partial evaluation of  $S$  in  $P$  using  $U$ .

<pre> while0 (X, Y, M) ←   X&gt;0,   if0 (X, Y, M) . while0 (X, Y, M) ←   X=&lt;0. </pre>	<pre> while0 (X, Y, M) ←   X&gt;0,   Y&lt;M, Y1=Y+1,   while0 (X, Y1, M) . while0 (X, Y, M) ←   X&gt;0,   Y&gt;=M, X1=X-1,   while0 (X1, Y, M) . while0 (X, Y, M) ←   X=&lt;0. </pre>	<pre> while0 (X, Y, M) ←   X&gt;0,   Y&lt;M, Y1=Y+1,   X&gt;0,   if0 (X, Y1, M) . while0 (X, Y, M) ←   X&gt;0,   Y&lt;M, Y1=Y+1,   X=&lt;0. while0 (X, Y, M) ←   X&gt;0,   Y&gt;=M, X1=X-1,   while0 (X1, Y, M) . while0 (X, Y, M) ←   X=&lt;0. </pre>
---	---	--

Figure 4: Three possible partial evaluations of constrained fact  $\text{while0}(X, Y, M) \leftarrow \text{true}$  in Figure 3(b). The leftmost column is the trivial unfolding, consisting of the original clauses for  $\text{while0}$ .

**Example 1** Let  $P$  be the set of clauses in Figure 3(b). Let  $A$  be the constrained atom  $\text{while0}(X, Y, M) \leftarrow \text{true}$ .

Figure 4 shows three sets of CHCs that are examples of partial evaluations of  $A$  in  $P$ , with different unfolding rules. Note that for CHCs with at most one atom in the body, such as in this example, the choice of atom for each derivation step is determinate, so the unfolding rule only determines how far to unfold each branch.

### 3 Specialisation algorithm

Algorithm 1 shows an outline specialisation algorithm SP for CHCs based on the “basic algorithm” in [10]. This is a so-called *online* partial evaluation algorithm, which makes control decisions on evaluation and polyvariance on the fly, as opposed to *offline* partial evaluation, in which control decisions are determined by the results of a prior analysis such as a binding time analysis.

The algorithm takes as input a set of CHCs  $P$  and a set of entry points  $S_0$ , where each entry point is a constrained fact. It is parameterised by two operations, namely  $\text{pe}$  and  $\alpha_p$ .

- $\text{pe}(S)$  returns a partial evaluation (Definition 5) of the set of constrained facts  $S$ .
- $\alpha_p(S)$  is a set of constrained facts such that  $\alpha_p(S) = \{p(\mathbf{x}) \leftarrow \rho(\phi) \mid p(\mathbf{x}) \leftarrow \phi \in S\}$ , where  $\rho$  is some generalisation operator (Definition 2).

Two other functions are called in the algorithm,  $\text{collect}$  and  $\text{rename}_\rho$ .

- $\text{collect}$ : Let  $Q$  be a set of CHCs.  $\text{collect}(Q)$  returns the set of constrained atoms collected from the bodies of clauses in  $Q$ . It is defined as  $\text{collect}(Q) = \{p_i(\mathbf{x}_i) \leftarrow \phi \mid p_0(\mathbf{x}_0) \leftarrow \phi, p_1(\mathbf{x}_1), \dots, p_k(\mathbf{x}_k) \in Q\}$ .

The definition of the  $\text{rename}_\rho$  function is postponed to Example 3.

The successive values of  $S'$  in the **repeat** loop of the algorithm (line 7) form an increasing sequence of sets of constrained facts with respect to  $\preceq$  starting from the input set  $S_0$ , say  $S_0, S_1, S_2, \dots$ ; the loop terminates if for some  $j > 0$ ,  $S_{j-1} = S_j$ .

**Algorithm 1**  $SP(P, S_0)$ 


---

```

1: Input: Finite set of CHCs  $P$ , finite set of constrained facts  $S_0$ , generalisation operator  $\rho$ .
2: Output: Finite set of CHCs
3:  $S \leftarrow S_0$ 
4: repeat
5:    $S' = S$ 
6:    $S \leftarrow S \cup \alpha_\rho(\text{collect}(\text{pe}(S)))$ 
7: until  $S' = S$ 
8: return  $\text{rename}_\rho(\text{pe}(S))$ 

```

---

## 4 Property-based abstraction

We now turn to the consideration of the abstraction function  $\alpha_\rho$ , using property-based abstraction.

We first define a generalisation operator  $\rho_\Psi$ . Let  $\Psi \subseteq \mathcal{C}^T$  be a finite set of constraints. Given a formula  $\phi \in \mathcal{C}^T$ , then

$$\rho_\Psi(\phi) = \bigwedge \{ \psi \mid \psi \in \Psi, \phi \preceq \psi \} \wedge \bigwedge \{ \neg \psi \mid \psi \in \Psi, \phi \preceq \neg \psi \}.$$

**Lemma 1**  $\rho_\Psi$  is a generalisation operator, that is, for all  $\phi \in \mathcal{C}^T$ ,  $\phi \preceq_T \rho_\Psi(\phi)$ .

Note that  $\rho_\Psi(\phi)$  is a conjunction of elements of  $\Psi$  and negations of elements of  $\Psi$ . Since we assume that  $\Psi$  is finite, then the set of possible values of  $\rho_\Psi(\phi)$  is finite. If none of the elements of  $\Psi$  or their negations is entailed by  $\phi$  then  $\rho_\Psi(\phi) = \text{true}$ .

**Example 2** Consider a set  $\Psi = \{ \psi_1, \psi_2, \psi_3 \}$ , where  $\psi_1 \wedge \psi_3 = \text{false}$ ,  $\psi_2 \wedge \psi_3 = \text{false}$  and  $\psi_1 \wedge \psi_2 \neq \text{false}$ . Figure 5 shows the generalisation for various choices of a property  $\phi$ , that is, the values of  $\rho_\Psi(\phi)$ .

$\rho_\Psi$  is extended to apply to constrained facts, and for convenience we take the set  $\Psi$  also to consist of constrained facts.

$$\rho_\Psi(p(\mathbf{x}) \leftarrow \phi) = p(\mathbf{x}) \leftarrow \bigwedge \{ \psi \mid p(\mathbf{x}) \leftarrow \psi \in \Psi, \phi \preceq \psi \} \wedge \bigwedge \{ \neg \psi \mid p(\mathbf{x}) \leftarrow \psi \in \Psi, \phi \preceq \neg \psi \}.$$

Let  $S$  and  $\Psi$  be sets of constrained facts. The operation  $\alpha_\rho$  from Algorithm 1 is defined where  $\rho$  is the generalisation operator  $\rho_\Psi$ .

$$\alpha_{\rho_\Psi}(S) = \{ p(\mathbf{x}) \leftarrow \rho_\Psi(\phi) \mid p(\mathbf{x}) \leftarrow \phi \in S \}$$

We now have all the components of the algorithm, and we show an example of specialisation with property-based abstraction.

**Example 3** Let  $P$  be the set of CHCs representing the code in Figure 2(a), that is, the clauses from Figure 3(b). Let  $S_0 = \{ \text{start} \leftarrow \text{true} \}$  and let the set  $\Psi$  contain the following constrained facts.

$$\begin{array}{lll} \text{while0}(A, B, C) \leftarrow A > 0 & \text{while0}(A, B, C) \leftarrow B < C & \text{if0}(A, B, C) \leftarrow B < C \\ \text{while0}(A, B, C) \leftarrow A \leq 0 & \text{while0}(A, B, C) \leftarrow B \geq C & \text{if0}(A, B, C) \leftarrow B \geq C \end{array}$$

The following unfolding rule is used: each branch is unfolded until either a branch point is reached (that is, a call to a predicate that appears in the head of more than one clause) or a recursive predicate is reached (more precisely, the predicate that is the target of a back edge in the predicate dependency graph of the program traversed from the initial predicate  $\text{start}$ ). In the given clauses, this implies that each partial evaluation consists of only one unfolding step. This unfolding rule guarantees termination of unfolding, but is conservative. We discuss other unfolding rules later. Algorithm 1 proceeds as follows.

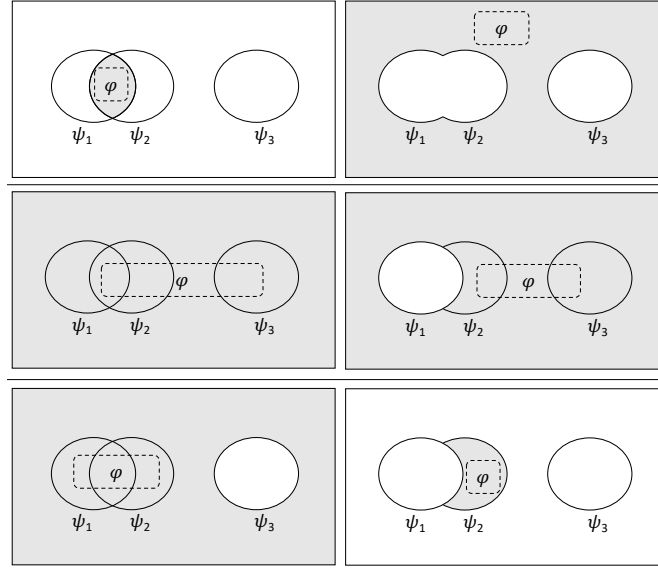


Figure 5: The property to be generalised ( $\varphi$ ) is shown as an area with dotted outline, and the shaded areas show the generalisation of  $\varphi$  using operation  $\rho_{\Psi}(\varphi)$ , where  $\Psi = \{\psi_1, \psi_2, \psi_3\}$ .

- *Initialisation:*  $S_0 = \{\text{start} \leftarrow \text{true}\}$ .
- *Iteration 1:*  $S_1 = S_0 \cup \{\text{while0}(A, B, C) \leftarrow \text{true}\}$ .
- *Iteration 2:*  $S_2 = S_1 \cup \{\text{if0}(A, B, C) \leftarrow A > 0\}$ .
- *Iteration 3:*  $S_3 = S_2 \cup \{\text{while0}(A, B, C) \leftarrow A > 0, \text{while0}(A, B, C) \leftarrow B \geq C\}$ .
- *Iteration 4:*  $S_4 = S_3 \cup \{\text{if0}(A, B, C) \leftarrow A > 0, B \geq C\}$ .
- *Iteration 5:*  $S_5 = S_4$ .

In this sequence we applied the operator  $\alpha_{\rho_{\Psi}}$  at each stage to reach the sets shown in the sequence. (To be precise, we apply the generalisation operator only to constrained facts for the recursive predicate `while0`, which is sufficient to ensure termination). For example, in iteration 3, partial evaluation of the constrained fact `if0(A, B, C) ← A > 0` results in the following CHCs.

$$\begin{aligned} \text{if0}(A, B, C) \leftarrow A > 0, C > B, D = B - 1, \text{while0}(A, D, C). \\ \text{if0}(A, B, C) \leftarrow A > 0, B \geq C, D = A - 1, \text{while0}(D, B, C). \end{aligned}$$

The constraints projected onto the body atoms `while0(A, D, C)` and `while0(D, B, C)` are respectively  $A > 0$  and  $D > -1, B \geq C$ . The result of applying abstraction, that is,

$$\alpha_{\rho_{\Psi}}(\{\text{while0}(A, D, C) \leftarrow A > 0, \text{while0}(D, B, C) \leftarrow D > -1, B \geq C\})$$

is  $\{\text{while0}(A, B, C) \leftarrow A > 0, \text{while0}(A, B, C) \leftarrow B \geq C\}$ . Notice that the second constrained fact has been generalised; evaluating  $\rho_{\Psi}(\text{while0}(D, B, C) \leftarrow D > -1, B \geq C)$  yields  $\text{while0}(A, B, C) \leftarrow B \geq C$ , which represents the only element of  $\Psi$  whose constraint is entailed. Also, in iteration 5, partial evaluation of `if0(A, B, C) ← A > 0, B ≥ C` results in the CHC:

$$\text{if0}(A, B, C) \leftarrow A > 0, B \geq C, D = A - 1, \text{while0}(D, B, C).$$




Figure 6: The predicate dependency graph for Example 3

The constraint on the body atom  $\text{while0}(D,B,C)$  is  $D > -1, B \geq C$ ; as before this is generalised to the constrained fact  $\text{while0}(A,B,C) \leftarrow B \geq C$ .

The function  $\text{rename}_p$  of Algorithm 1 gives each constrained atom in the final set a different name, in this case as follows.

Constrained fact	VersionName	Constrained fact	VersionName
$\text{start} \leftarrow \text{true}$	start	$\text{if0}(A,B,C) \leftarrow A > 0, B \geq C$	$\text{if0}_1(A,B,C)$
$\text{while0}(A,B,C) \leftarrow B \geq C$	$\text{while0}_2(A,B,C)$	$\text{while0}(A,B,C) \leftarrow A > 0$	$\text{while0}_3(A,B,C)$
$\text{if0}(A,B,C) \leftarrow A > 0$	$\text{if0}_4(A,B,C)$	$\text{while0}(A,B,C) \leftarrow \text{true}$	$\text{while0}_5(A,B,C)$

Finally, the partial evaluations of all the constrained atoms in the final set are collected, and the atoms renamed according to the versions that occur in the final set. This gives the following specialised CHC clauses.

```

start ← while05(A,B,C).
while05(A,B,C) ← A>0, if04(A,B,C).
while05(A,B,C) ← -A>=0.
if04(A,B,C) ← A>0, -B+C>0, B+ -D= -1, while03(A,D,C).
if04(A,B,C) ← A>0, B+ -C>=0, A+ -D=1, while02(D,B,C).
while03(A,B,C) ← A>0, if04(A,B,C).
while02(A,B,C) ← B+ -C>=0, A>0, if01(A,B,C).
while02(A,B,C) ← B+ -C>=0, -A>=0.
if01(A,B,C) ← A>0, B+ -C>=0, A+ -D=1, while02(D,B,C).

```

The predicate dependency graph for these clauses is shown in Figure 6 and it can be seen that this has the same structure as the code in Figure 2(b) where there are two distinct loops. Note that the predicate  $\text{while0}_5$  is in fact not the head of a loop but rather the initial if-statement of Figure 2(b), while the predicate  $\text{if0}_4$  is the head of a loop.

## 5 Choice of properties and granularity of abstraction

The set of properties  $\Psi$  in Example 3 was chosen so that the properties were relevant to the tests determining the control flow. However, the choice of properties can be critical to achieving good specialisations. In this section we discuss the effect of the choice of properties.



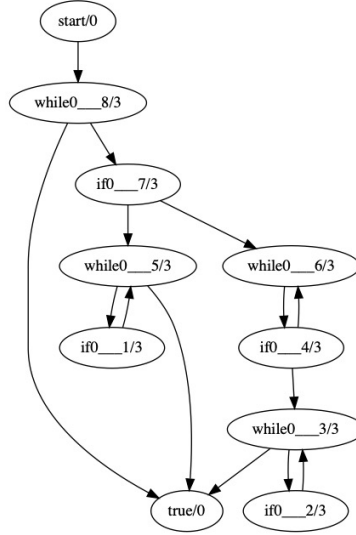


Figure 7: The predicate dependency graph for Example 3, with enlarged set of properties.

In general, it is clear that the larger the set of properties, the more versions of predicates can be produced, and thus more specialised clauses can be generated. Fewer properties, on the other hand, cause information needed for specialisation to be lost. For example, the following properties could also be chosen in Example 3. It is a subset of the set previously chosen, incorporating only the constraints directly appearing in the clauses for the respective predicates.

$$\{ \text{while0}(A, B, C) \leftarrow A > 0, \text{while0}(A, B, C) \leftarrow A \leq 0, \\ \text{if0}(A, B, C) \leftarrow B < C, \text{if0}(A, B, C) \leftarrow B \geq C \}$$

Using this choice of  $\Psi$ , no specialisation at all is achieved; the original clauses are returned. The problem is that the constraints on `if0` are lost when abstracting the calls to `while0`, since the properties applying to `while0(A, B, C)` say nothing about the values of `B` or `C`.

However, for a given unfolding rule, there is a limit to how much specialisation can be achieved, no matter how many properties  $\Psi$  contains. Consider the following set  $\Psi$  for Example 3, which results from collecting all constraints from the given clauses, projected onto head and body atoms.

$$\begin{array}{ll} \text{while0}(A, B, C) \leftarrow A > 0 & \text{if0}(A, B, C) \leftarrow A > 0 \\ \text{while0}(A, B, C) \leftarrow A \leq 0 & \text{if0}(A, B, C) \leftarrow B < C \\ \text{while0}(A, B, C) \leftarrow C > B - 1 & \text{if0}(A, B, C) \leftarrow B \geq C \\ \text{while0}(A, B, C) \leftarrow B \geq C & \end{array}$$

Figure 7 shows the predicate dependency graph for the clauses resulting from this set. While there are more versions of predicates than in Figure 6, the corresponding clauses are no more specialised. Viewing the graph as a finite automaton, it can be verified that the states `if0_4` and `if0_7` are equivalent, as are `while0_5` and `while0_3`, and `if0_1` and `if0_2`. In short, the automaton can be minimised to give the same automaton in Figure 6. More specifically, the versions corresponding to constrained facts  $\text{if0}(A, B, C) \leftarrow A > 0, C > B - 1$  and  $\text{if0}(A, B, C) \leftarrow A > 0$  are indistinguishable using the given unfolding rule. Intuitively, there can be no better specialisation for the input clauses in Example 3, with any set of properties  $\Psi$  or indeed any other unfolding rule, than the one achieved in Example 3.

Space does not permit a detailed account of the automata-theoretic argument, but we state the fol-

<pre> int a, b; if (a ≤ 100)   a = 100-a; else   a=a-100; while (a ≥ 1)   a=a-1;   b=b-2; assert (b != 0); </pre>	<pre> init(A,B) ← true. if(A,B) ← A0 ≤ 100, A=100-A0, init(A0,B). if(A,B) ← A0 ≥ 101, A=A0-100, init(A0,B). while(A,B) ← if(A,B). while(A,B) ← A0 ≥ 1, A=A0-1, B=B0-2,                while(A0,B0). false ← A ≤ 0, B=0,                while(A,B) </pre>
---	--

Figure 8: Example from [16]: (left) original program, (right) translation to CHCs

lowing proposition.

**Proposition 1** *Let  $P$  be a set of CHCs, specialised to a set  $P'$  using Algorithm 1 using a given input set  $S_0$ , set of properties  $\Psi$  and unfolding rule. Then there exists a set of clauses  $P''$  having a minimal number of versions of each predicate, that is equivalent to  $P'$  with respect to computations starting from  $S_0$ .*

We also conjecture that there exists a set of properties  $\Psi'$  such that executing the algorithm with  $\Psi'$  would yield this minimal version.

Proposition 1 holds for a fixed unfolding rule. Of course, further specialisation might be achievable using a more aggressive unfolding rule.

## 6 Polyvariant specialisation in verification

In previous work, we have used specialisation of CHCs as a component in program verification tasks. In this section we discuss its role.

**Pre-condition inference.** Firstly, in [16], Algorithm 1 was used as a component in an algorithm for computing sufficient conditions for safety of imperative programs encoded as CHCs. In many cases, the required safety condition is a disjunction. Polyvariant specialisation enabled the relevant disjuncts to be found by a convex polyhedral analysis.

**Example 4** *Consider the example in Figure 8 taken from [16]. Note that the translation to CHCs corresponds to a backwards flow of control from the error predicate `false` to the program start predicate `init`. The goal is to infer conditions on the start predicate that ensure that the error predicate is not reached. Specialisation of the set of CHCs was carried out using the following set of properties, and the same unfolding rule as in Example 3.*

```

{if(A,B) ← A ≥ 0, if(A,B) ← A ≥ 1,
init(A,B) ← A ≥ 101, init(A,B) ← -A ≥ -100,
while(A,B) ← A ≥ 0, while(A,B) ← A ≥ 1, while(A,B) ← -A ≥ 0, B=0,
while(A,B) ← -A ≥ 0, while(A,B) ← B=0}

```

*It can be seen that three versions of the predicate `init` have been generated, arising from different paths through the program. Analysis of these specialised CHCs allowed the disjunctive precondition on `init(A,B)`, namely  $B \neq |2A - 100|$  to be derived. This condition could not be derived from the original code without an analysis domain of disjunctive properties, which is more difficult to implement and control. Polyvariant specialisation, in effect, provides a heuristic for introducing disjunctions selectively, where they can affect the control flow.*

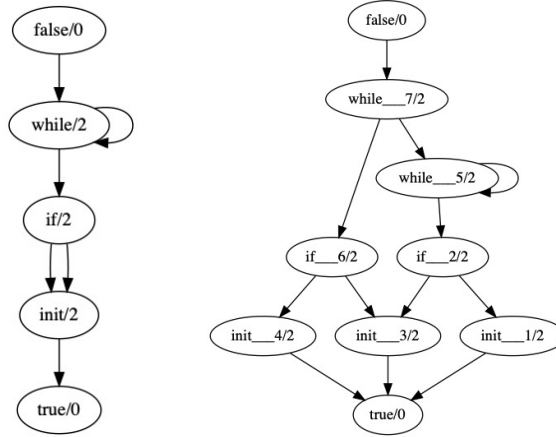


Figure 9: The predicate dependency graph for Example 4, before and after polyvariant specialisation

**Termination analysis.** In [8], polyvariant specialisation was used to transform a control-flow graph obtained from a program into another equivalent control-flow graph in which the loops were in a form more suitable for automatic proof of termination. The control-flow graphs are straightforwardly represented as CHCs and examples of their polyvariant specialisation are displayed in [8]. The example from Figure 2 provides a case in point. The structure of the single loop makes it rather hard to find a suitable ranking function that establishes termination; whereas the restructured code based on polyvariant specialisation, with two separate loops, is easy to prove terminating, since each loop has a simple ranking function. A large number of experiments was reported in [8], showing that polyvariant specialisation of the control-flow graph very often improves the effectiveness of both automatic termination analysis and complexity bound analysis. In short, the work demonstrates that control-flow refinement [14] can be achieved by polyvariant specialisation. Another relevant approach for finding better loop invariants, is the “splitter predicates” method [23]; this can also be reproduced using property-based specialisation.

**Dimension-based decomposition.** The concept of *tree dimension* has been applied in verification to decompose a proof. The dimension of a set of CHCs is a measure of their non-linearity. A set of CHCs of dimension zero contains only linear clauses (that is, clauses having at most one atom in the body). Proof trees in such sets of CHCs have no branching. Sets of clauses of higher dimension give rise to branching proof trees, and the dimension of a tree is determined by the dimensions of the subtrees of the root, as illustrated in Figure 10. A more detailed definition can be found in [17]. Given a set of CHCs  $P$  for which some property is to be verified, we may decompose the problem by dimensions. For a given dimension  $d$ , we can define a set of CHCs, say  $P^d$ , such that an atom  $A$  has a proof in  $P$  of dimension  $d$  if and only if it has a proof in  $P^d$ . Since every proof has some finite dimension,  $P \vdash A$  if and only if  $P^0 \vdash A \vee P^1 \vdash A \vee P^2 \vdash A \dots$ . In [17], we showed a technique using polyvariant specialisation for generating various dimension-bounded sets of clauses. We could generate  $P^d$ , the set of clauses yielding exactly the proof trees of dimension  $d$ ; we could also generate  $P^{\leq d}$ , the set of clauses yielding proof trees of dimension at most  $d$ ; and we could generate  $P^{> d}$ , the set of clauses yielding proof trees of dimension at least  $d + 1$ . This enabled a variety of strategies for proof decomposition, described in detail in [17].

In order to derive such dimension-constrained sets of clauses, we first instrumented the clauses with an extra argument for each predicate, representing the dimension, together with constraints capturing the




Figure 10: The dimension of a node in a tree: leaf nodes have dimension 0; a node has dimension  $n + 1$  if at least two subtrees have dimension  $n$ ; it has dimension  $n$  if exactly one subtree has dimension  $n$  and any other subtrees have lower dimension.

rule for computing dimension. That is, the clause  $p_0(\mathbf{x}_0) \leftarrow \phi, p_1(\mathbf{x}_1), \dots, p_n(\mathbf{x}_n)$  is replaced by

$$p_0(\mathbf{x}_0, k) \leftarrow \phi, p_1(\mathbf{x}_1, k_1), \dots, p_n(\mathbf{x}_n, k_n), \text{dim}(k_1, \dots, k_n, k)$$

where  $\text{dim}(k_1, \dots, k_n, k)$  represents the computation of the head dimension  $k$  from the (subtree) dimensions  $k_1, \dots, k_n$ .

Specialisation could then be performed by providing constraints on the argument  $k$ . The set of properties  $\Psi$  input to Algorithm 1 consisted a constrained fact for each dimension up to the required bound.

**Example 5** Let  $P$  be the set of clauses for the Fibonacci function, instrumented with the dimension as described above, together with a constraint representing a property to be proved and a dimension bound on false of 2.

```

fib(A, B, 0) :- A >= 0, A <= 1, A = B.
fib(A, B, K) :- A > 1, D = A - 2, E = A - 1, B = F + G, fib(D, G, K2), fib(E, F, K1),
               K1 + 1 = < K, K2 = K.
fib(A, B, K) :- A > 1, D = A - 2, E = A - 1, B = F + G, fib(D, G, K1), fib(E, F, K2),
               K1 + 1 = < K, K = K2.
fib(A, B, K) :- A > 1, D = A - 2, E = A - 1, B = F + G, fib(D, G, K1), fib(E, F, K2),
               K1 = K - 1, K2 = K1.
false(A) <- X > 5, fib(X, Y, K), Y < X, K <= 2.

```

Let  $\Psi$  be the following set of constrained facts.

```

fib(A, B, C) <- C <= 2, fib(A, B, C) <- C <= 1, fib(A, B, C) <- C <= 0, fib(A, B, C) <- C >= 0,
false(A) <- A <= 2, false(A) <- A <= 1, false(A) <- A <= 0, false(A) <- A >= 0,

```

Figure 11 shows the predicate dependency graphs before and after polyvariant specialisation using initial call  $\text{false}(A) \leftarrow A \leq 2$ ,  $\Psi$  shown above and an unfolding rule that just unfolds one step. Here,  $\text{fib}_1$ ,  $\text{fib}_2$  and  $\text{fib}_3$  yield proof trees of dimension  $\leq 0$ ,  $\leq 1$  and  $\leq 2$  respectively.

## 7 Discussion and related work

The control of partial evaluation, and more broadly of program specialisation, has been much studied [15]. The problem arises due to the two termination problems in specialisation: local termination, ensur-

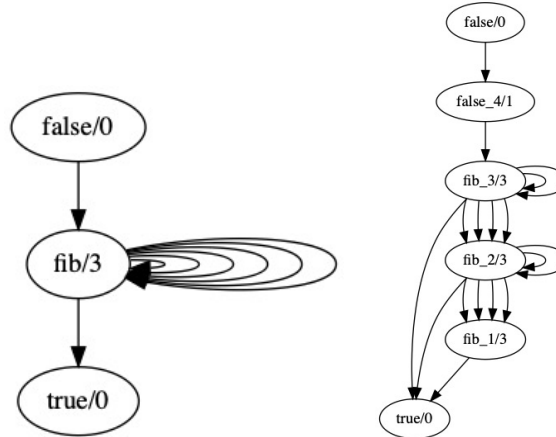


Figure 11: The predicate dependency graph for Example 5 producing clauses yielding proof trees of dimension at most 2, before and after polyvariant specialisation

ing that loops in the program are not unfolded indefinitely, and global termination, which can be seen as the problem of generating only a finite number of versions of each program point. These two problems are mutually dependent: the more conservative the local unfolding strategy is (in order to ensure local termination), the more important it is to allow multiple versions of program points in order to preserve information, thus increasing the risk of global non-termination. In some approaches, the two problems are somewhat merged in order to handle these interactions [24, 22, 21, 18]. However, from the point of view of conceptual clarity and implementability, we argue that it is desirable to separate the two problems. Thus it becomes important to have a flexible way of controlling polyvariance and global termination.

The problem of polyvariant specialisation also occurs in offline partial evaluation; the unfolding of a program point marked as static by a binding time analysis must be accompanied by an assurance that only a finite number of static instances will arise at partial evaluation time, and thus a finite number of versions of the program point will be generated. This is known as bounded static variation [15].

The algorithm presented here is based on a framework for partial evaluation of logic programs originally formulated in [2] and refined in [10], based on the framework presented by Lloyd and Shepherdson [20]. Some of the definitions in Section 2 are from these works, adapted for constrained Horn clauses.

The concept of property-based abstraction has been widely used in software model checking and was first introduced by Ball *et al.* [1] where it is called the Cartesian abstraction. It is used in a form similar to that shown in this paper in the HSF tool [13], though that work does not use the negations of the properties as we do (following [1]). Though in this work we do not present it as an abstract interpretation [4], the domain of properties based on a finite set of constraints forms a lattice and the process of abstracting a concrete property is an example of a Galois connection. Being a finite domain, it has both advantages and disadvantages compared to other domains that could be used to control polyvariance. For a fixed local unfolding rule, property-based abstraction places a fixed limit on the number of realisable versions, whereas with an infinite-height abstract domain, with global termination ensured by widening, the dependence of global termination on local unfolding would be loosened and an infinite number of versions could be produced. A general presentation of the relation between fixed height versus infinite height domains can be found in [5], while an attempt to implement global control using an infinite height lattice is shown in [12].

Property-based abstractions are indirectly related to trace-based abstractions, which have been used in partial evaluation and supercompilation to control polyvariance, e.g. [24, 18, 11, 6]. Properties determine traces and vice versa; a property constrains the feasible program traces; whereas a trace implicitly defines properties which permit the trace (a principle explicitly used by de Angelis *et al.* [6]). We consider property-based abstractions to have some practical and conceptual advantages, opening up the use of satisfiability solvers to compute the abstraction. However, further evaluation is needed, as well as investigation on the choice of properties.

The usefulness of specialisation as a component in program verification tools has been established in many works, including [25, 19, 7, 16] to name only a few. Fioravanti *et al.* investigated the trade-offs of polyvariance with efficiency and precision when using specialisation as a verification tool [9]. The use of constrained Horn clauses as a semantic representation formalism for verification of a wide range of languages and systems is now well established [13, 3]. Here, we emphasise the role of polyvariant specialisation in generating multiple versions of program points in a controlled way, when such versions lead to different control flow. This in turn implicitly allows disjunctive properties to be handled.

**Acknowledgements.** The author acknowledges very useful discussions with Bishoksan Kafle, Pierre Ganty, Samir Genaim and Jesús Doménech, contributing to the ideas presented here.

## References

- [1] T. Ball, A. Podelski & S. K. Rajamani (2001): *Boolean and Cartesian Abstraction for Model Checking C Programs*. In T. Margaria & W. Yi, editors: *Tools and Algorithms for the Construction and Analysis of Systems, 7th International Conference, TACAS 2001, Proceedings, Lecture Notes in Computer Science 2031*, Springer, pp. 268–283, doi:10.1007/3-540-45319-9\_19.
- [2] K. Benkerimi & J. W. Lloyd (1990): *A Partial Evaluation Procedure for Logic Programs*. In S. K. Debray & M. V. Hermenegildo, editors: *Logic Programming, Proceedings of the 1990 North American Conference*, MIT Press, pp. 343–358.
- [3] N. Bjørner, A. Gurfinkel, K. L. McMillan & A. Rybalchenko (2015): *Horn Clause Solvers for Program Verification*. In L. D. Beklemishev, A. Blass, N. Dershowitz, B. Finkbeiner & W. Schulte, editors: *Fields of Logic and Computation II - Essays Dedicated to Yuri Gurevich on the Occasion of His 75th Birthday, Lecture Notes in Computer Science 9300*, Springer, pp. 24–51, doi:10.1007/978-3-319-23534-9\_2.
- [4] P. Cousot & R. Cousot (1977): *Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints*. In: *Proceedings of the 4th ACM Symposium on Principles of Programming Languages, Los Angeles*, pp. 238–252.
- [5] P. Cousot & R. Cousot (1992): *Comparing the Galois Connection and Widening/Narrowing Approaches to Abstract Interpretation*. In: *Proceedings of the 4th International Symposium on Programming Language Implementation and Logic Programming, Springer-Verlag Lecture Notes in Computer Science 631*, pp. 269–295.
- [6] E. De Angelis, F. Fioravanti, A. Pettorossi & M. Proietti (2012): *Specialization with Constrained Generalization for Software Model Checking*. In E. Albert, editor: *Logic-Based Program Synthesis and Transformation, LOPSTR 2012, Lecture Notes in Computer Science 7844*, Springer, pp. 51–70, doi:10.1007/978-3-642-38197-3\_5.
- [7] E. De Angelis, F. Fioravanti, A. Pettorossi & M. Proietti (2014): *Program verification via iterated specialization*. *Sci. Comput. Program.* 95, pp. 149–175.
- [8] J. Doménech, S. Genaim & J. P. Gallagher (2018): *Control-Flow Refinement via Partial Evaluation*. In S. Lucas, editor: *16th International Workshop on Termination (WST 2018)*, pp. 55–59. Available at <http://wst2018.webs.upv.es/wst2018proceedings.pdf>.

- [9] F. Fioravanti, A. Pettorossi, M. Proietti & V. Senni (2013): *Controlling Polyvariance for Specialization-based Verification*. *Fundam. Inform.* 124(4), pp. 483–502, doi:10.3233/FI-2013-845.
- [10] J. P. Gallagher (1993): *Specialisation of Logic Programs: A Tutorial*. In: *Proceedings PEPM'93, ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, ACM Press, Copenhagen, pp. 88–98.
- [11] J. P. Gallagher & L. Lafave (1996): *Regular Approximation of Computation Paths in Logic and Functional Languages*. In O. Danvy, R. Glück & P. Thiemann, editors: *Partial Evaluation, Springer-Verlag Lecture Notes in Computer Science* 1110, pp. 115–136.
- [12] J. P. Gallagher & J. C. Peralta (2001): *Regular Tree Languages as an Abstract Domain in Program Specialisation*. *Higher-Order and Symbolic Computation* 14(2-3), pp. 143–172.
- [13] S. Grebenschikov, N. P. Lopes, C. Popeea & A. Rybalchenko (2012): *Synthesizing software verifiers from proof rules*. In J. Vitek, H. Lin & F. Tip, editors: *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12*, ACM, pp. 405–416.
- [14] S. Gulwani, S. Jain & E. Koskinen (2009): *Control-flow refinement and progress invariants for bound analysis*. In M. Hind & A. Diwan, editors: *PLDI 2009*, ACM, pp. 375–385, doi:10.1145/1542476.1542518.
- [15] N. D. Jones, C. Gomard & P. Sestoft (1993): *Partial Evaluation and Automatic Software Generation*. Prentice Hall.
- [16] B. Kafle, J. P. Gallagher, G. Gange, P. Schachte, H. Søndergaard & P. J. Stuckey (2018): *An iterative approach to precondition inference using constrained Horn clauses*. *TPLP* 18(3-4), pp. 553–570.
- [17] B. Kafle, J. P. Gallagher & P. Ganty (2018): *Tree dimension in verification of constrained Horn clauses*. *TPLP* 18(2), pp. 224–251.
- [18] M. Leuschel & B. Martens (1996): *Global Control for Partial Deduction through Characteristic Atoms and Global Trees*. In O. Danvy, R. Glück & P. Thiemann, editors: *Partial Evaluation, Springer-Verlag Lecture Notes in Computer Science* 1110, pp. 263–283.
- [19] M. Leuschel & T. Massart (2000): *Infinite State Model Checking by Abstract Interpretation and Program Specialisation*. In A. Bossi, editor: *Logic-Based Program Synthesis and Transformation (LOPSTR'99)*, Springer-Verlag Lecture Notes in Computer Science 1817, pp. 63–82.
- [20] J. Lloyd & J. Shepherdson (1991): *Partial Evaluation in Logic Programming*. *Journal of Logic Programming* 11(3 & 4), pp. 217–242.
- [21] B. Martens & J. P. Gallagher (1995): *Ensuring Global Termination of Partial Deduction While Allowing Flexible Polyvariance*. In L. Sterling, editor: *Proc. International Conference on Logic Programming, (ICLP'95)*, Tokyo, MIT Press.
- [22] D. Sahlin (1993): *Mixtus: An Automatic Partial Avaluator for Full Prolog*. *New Generation Comput.* 12(1), pp. 7–51, doi:10.1007/BF03038271.
- [23] R. Sharma, I. Dillig, T. Dillig & A. Aiken (2011): *Simplifying Loop Invariant Generation Using Splitter Predicates*. In G. Gopalakrishnan & S. Qadeer, editors: *Computer Aided Verification, CAV 2011, Lecture Notes in Computer Science* 6806, Springer, pp. 703–719, doi:10.1007/978-3-642-22110-1\_57.
- [24] V. Turchin (1988): *The Algorithm of generalization in the supercompiler*. In D. Bjørner, A. Ershov & N. Jones, editors: *Proc. of the IFIP TC2 Workshop on Partial Evaluation and Mixed Computation*, North-Holland, pp. 531–549.
- [25] D. de Waal & J. P. Gallagher (1994): *The Applicability of Logic Program Analysis and Transformation to Theorem Proving*. In: *Proceedings of the 12th International Conference on Automated Deduction (CADE-12)*, Nancy.