

Satisfiability checking for PC(ID)

Maarten Mariën Rudradeb Mitra
Marc Denecker Maurice Bruynooghe

Report CW 426, October 2005



Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Satisfiability checking for PC(ID)*

Maarten Mariën Rudradeb Mitra
Marc Denecker Maurice Bruynooghe

Report CW426, October 2005

Department of Computer Science, K.U.Leuven

Abstract

The logic FO(ID) extends classical first order logic with inductive definitions. This paper studies the satisfiability problem for PC(ID), its propositional fragment. We develop a framework for model generation in this logic, present an algorithm and prove its correctness. As FO(ID) is an integration of classical logic and logic programming, our algorithm integrates techniques from SAT and ASP. We report on a prototype system, called MIDL, experimentally validating our approach.

AMS(MOS) Classification : Primary : 68T30, Secondary : 68T27.

*This research was supported by FWO-Vlaanderen, IWT-Vlaanderen, European Framework 5 Project WASP, and by GOA/2003/08.

1 Introduction

The logic FO(ID), or Inductive Definition Logic (ID-logic) [7], extends classical first order logic (FO) with a language primitive that allows a uniform representation of inductive definitions. In general, inductive definitions cannot be represented in first order logic (FO). The semantics of this primitive is based on the well-founded semantics of logic programming [29]; indeed, as argued in [6, 8], it correctly formalizes the semantics of inductive definitions.

While definitions are common in mathematics, they are also crucial in declarative Knowledge Representation. Not only non-inductive definitions are frequent in common-sense reasoning as argued in the seminal paper [2], also inductive definitions are. In [10], the *situation calculus* is given a very natural and general representation as an iterated inductive definition in the well-ordered set of situations and [17] observes that inductive definitions are present in many applications of Answer Set Programming (ASP) [12]. In short, definitions are a distinctive and important form of knowledge that can be naturally represented in FO(ID).

The goal of this paper is to present algorithms to solve SAT(PC(ID)), the satisfiability problem for PC(ID)¹, the propositional fragment of FO(ID), or equivalently, generate models for theories in this fragment. This problem is an extension of SAT, the satisfiability problem of propositional CNF formulas. Current SAT solvers exhibit impressive performance on many industrial instances. Unfortunately, SAT is a rather poor modeling language, and substantial effort is often required to encode a problem. PC(ID) is a major enhancement of the expressivity [19]. Solvers for SAT(PC(ID)) are also strongly related to ASP solvers such as Smodels [21] and DLV [4]. These solvers use the fixpoint operator of the well-founded semantics as a boolean propagation mechanism for rule sets.

Viable approaches for building a solver for SAT(PC(ID)) are:

1. A native approach that integrates inference techniques for inductive definitions with SAT inference techniques.
2. Mapping PC(ID) theories to CNF theories and applying off-the-shelf SAT solvers as in [23].
3. Mapping PC(ID) theories to equivalent general logic programs and then applying off-the-shelf ASP solvers. Using results from [17], it is easy to accomplish, but it completely bypasses PC(ID)'s relation with SAT.

The first approach, explored in this paper, is the most promising in the long run. It will improve our understanding of how to compute with inductive definitions and is best suited to integrate language extensions such as aggregates, constraints and open functions.

Despite the semantical differences, the computational tasks of our algorithm are very similar to those of algorithms like Smodels' Expand [21] and DLV's DetCons [4]. Its novelty lies in the use of justification semantics [9], which offers a different view on the computational task involved:

- One can locally test whether the well-founded operator is required.

¹ PC(ID) stands for “Propositional Calculus, extended with Inductive Definitions”.

- One can use a watched literal technique for propagations, very similar to the Two Watched Literal technique used in SAT [20].

We introduce PC(ID) in Section 2 and the semantic foundations of our algorithm in Section 3. We present the algorithm and argue its correctness in Section 4. Its implementation, MIDL, and a comparison with other approaches are described in Section 5. We finish with conclusions and related work.

2 Preliminaries

2.1 PC(ID)

The semantics of FO(ID) ([7]) is here redefined for the propositional fragment PC(ID). A vocabulary Σ is a set of atom symbols. A definition D is a set of rules of the form $(P \leftarrow \varphi)$, where $P \in \Sigma$ is called the head of the rule, and the body of the rule, φ , is a propositional formula in Σ . We denote by $Def(D)$ the set of atoms that appear in the heads of the rules of D . The set $\Sigma \setminus Def(D)$ is called the set of open symbols of D and is denoted by $Open(D)$. A literal is an atom P or its negation $\neg P$; it is *defined* if $P \in Def(D)$, otherwise it is *open*. A PC(ID) theory T in Σ is a set of propositional formulas and definitions in Σ .

A three-valued Σ -interpretation I is a function from Σ to the set $\{\mathbf{f}, \mathbf{u}, \mathbf{t}\}$ of truth values; a two-valued interpretation maps to $\{\mathbf{f}, \mathbf{t}\}$ instead. Truth values are ordered by the truth order, defined as $\mathbf{f} < \mathbf{u} < \mathbf{t}$, and the precision order, given by $\mathbf{u} <_p \mathbf{t}$ and $\mathbf{u} <_p \mathbf{f}$. We also have $\mathbf{f}^{-1} = \mathbf{t}$, $\mathbf{u}^{-1} = \mathbf{u}$ and $\mathbf{t}^{-1} = \mathbf{f}$. We denote the projection of I on the atoms in a set \mathcal{S} by $I|_{\mathcal{S}}$.

The semantics can be defined by means of the well-founded semantics [29]: Given a definition D and an interpretation $I|_{Open(D)}$ of the open atoms of D , there is a unique well-founded model of D extending $I|_{Open(D)}$, which we denote by $wfm_D(I|_{Open(D)})$. An interpretation I is a model of definition D , denoted $I \models D$, iff I is two-valued and $I = wfm_D(I|_{Open(D)})$; I is a model of a PC(ID) theory T iff I is two-valued and is a model of every definition and every propositional formula in T . An equivalent characterisation is provided by Corollary 1 below.

Example 1. Consider the definition $D_1 = \{P \leftarrow Q\}$. Then $P \in Def(D_1)$, $Q \in Open(D_1)$. The models of D_1 are $\{P, Q\}$ and \emptyset . They are also the models of $D_2 = \{Q \leftarrow P\}$. Hence, $T_1 = \{\{P \leftarrow Q\}, \{Q \leftarrow P\}\}$ which consists of two definitions has models $\{P, Q\}$ and \emptyset . However, $T_2 = \{\{P \leftarrow Q, Q \leftarrow P\}\}$ consists of a single definition and has only \emptyset as model.

The theory $T_3 = \{\{P \leftarrow Q\}, \{P \leftarrow \}, \neg Q \vee R\}$ has two definitions for P and one propositional formula; the first definition has two models (as D_1), while the second one has only $\{P\}$ as its model and $\{P, Q, R\}$ is the only model of T_3 .

2.2 MIDL normal form

Definition 1 (MIDL normal form). A PC(ID) theory T is in MIDL normal form iff $T = C \cup \{D\} \cup E$ where C is a set of clauses without defined literals,

E is a set of equivalences $P \equiv Q$ where P is an open and Q a defined atom, and D is a definition (a set of rules), with for each atom Q in $Def(D)$ exactly one rule with Q in the head. Moreover for all rules R in D , R is in the form $Q \leftarrow L_1 \wedge \dots \wedge L_n$ or $Q \leftarrow L_1 \vee \dots \vee L_n$, where $n \geq 1$ and L_i are literals.

Example 2. $T_3^{MidL} = \{\{P^{DA} \leftarrow Q, P^{DB} \leftarrow \}, \neg Q \vee R, P \equiv P^{DA}, P \equiv P^{DB}\}$ is the MIDL normal form equivalent to T_3 in Example 1: the models of T_3^{MidL} , restricted to T_3 's vocabulary, are the models of T_3 .

Defined atoms whose body is a disjunction respectively conjunction are called *disjunctive* respectively *conjunctive* atoms. Let \mathcal{S} be a set of literals; then we abbreviate $\bigwedge_{L \in \mathcal{S}} L$ by $\bigwedge \mathcal{S}$ and $\bigvee_{L \in \mathcal{S}} L$ by $\bigvee \mathcal{S}$.

The following transformation maps an arbitrary PC(ID) theory T to a theory in MIDL normal form. The transformation is linear in the size of the input theory, and is faithful, i.e. the input theory and the resulting output theory have the same models.

- for each sentence $\varphi \in T$, apply Tseitin's encoding [27] to transform it to clausal form
- for each rule in each definition: if its body is not a conjunction of literals, apply Lloyd-Topor's [15] transformation
- for each definition $D \in T$ and each $P \in Def(D)$: introduce a new atom P^D , and add the equivalence $(P \equiv P^D)$ to E . We call P the *original* of P^D ;
- replace each definition $D \in T$ with a new definition, which is obtained from D by replacing any occurrence of $P \in Def(D)$ by P^D ;
- if a rule defining P^D has empty body, replace the rule by $P^D \leftarrow t$ and add the unit clause t in C , where t is a new atom;
- if $P^D \leftarrow \varphi_1, \dots, P^D \leftarrow \varphi_n$ are the rules defining P^D and $n > 1$, replace these by the rules $P^D \leftarrow P_{\varphi_1}^D \vee \dots \vee P_{\varphi_n}^D, P_{\varphi_1}^D \leftarrow \varphi_1, \dots, P_{\varphi_n}^D \leftarrow \varphi_n$, where $P_{\varphi_i}^D$ are new atoms²;
- merge all definitions.

As the above example and the transformation show, the head of a definition $P \leftarrow \dots$ becomes a defined atom P^D which is linked to the original atom P through an equivalence.

3 Semantic Background

This section introduces semantical concepts which lie at the basis of our inference technique for inductive definitions, and are needed to prove correctness of our algorithms. These concepts were originally defined in [9].

Definition 2 (Direct justification). *Let D be a definition in MIDL normal form and J_d a set of literals. J_d is a direct justification for a defined atom P iff:*

² If φ_i is just one literal, it is not necessary to introduce a new atom and rule for it.

- either $P \leftarrow L_1 \wedge \dots \wedge L_n \in D$ and $J_d = \{L_1, \dots, L_n\}$;
- or $P \leftarrow L_1 \vee \dots \vee L_n \in D$ and $J_d = \{L_i\}$ for some $i \in 1 \dots n$.

J_d is a direct justification for a defined literal $\neg P$ iff:

- either $P \leftarrow L_1 \wedge \dots \wedge L_n \in D$ and $J_d = \{\neg L_i\}$ for some $i \in 1 \dots n$;
- or $P \leftarrow L_1 \vee \dots \vee L_n \in D$ and $J_d = \{\neg L_1, \dots, \neg L_n\}$.

Example 3. Consider $D_1 = \{P \leftarrow Q \wedge \neg R, Q \leftarrow \neg P \vee R\}$. The set $\{Q, \neg R\}$ is a direct justification for P and $\{P, \neg R\}$ for $\neg Q$; both $\{\neg P\}$ and $\{R\}$ are direct justifications for Q and both $\{\neg Q\}$ and $\{R\}$ for $\neg P$.

Consider $D_2 = \{P \leftarrow Q, Q \leftarrow P \vee R, R \leftarrow \neg P \wedge S\}$. Though $\{Q\}$ is a direct justification for P and $\{P\}$ for Q , the truth of P doesn't justify the truth of Q or vice versa, as illustrated by the model $\text{wfm}_D(\{S^f\}) = \{P^f, Q^f, R^f, S^f\}$.

A direct justification is insufficient to infer the truth of a literal in the well-founded model. We need to consider graphs of direct justifications. A *leaf* of a graph is a node without outgoing edges.

Definition 3 (Justification). A justification J of a definition in MIDL normal form is a directed graph where the nodes are literals, such that for each internal node L , L is a defined literal and the set of its children, $\text{Ch}_J(L)$, is a direct justification for L . A justification is total if none of its leaves are defined literals.

Example 4. D_2 from Example 3 has many justifications. Examples are:

$$J_0 = \emptyset, J_1 = \begin{array}{c} \boxed{P} \\ \uparrow \downarrow \\ \boxed{Q} \end{array}, J_2 = \begin{array}{ccccc} P & \longrightarrow & Q & \longrightarrow & R & \longrightarrow & S \\ \uparrow & & & & \downarrow & & \\ -R & \longleftarrow & -Q & \rightleftharpoons & -P & & \end{array}, J_3 = \begin{array}{c} \boxed{\neg P} \\ \uparrow \downarrow \\ \boxed{-Q} \longrightarrow \boxed{-R} \end{array}.$$

Here J_0 , J_1 and J_2 are total, but J_3 is not since $\neg R$ is a defined leaf.

As the example shows, justifications can contain cycles. We distinguish between *positive*, *negative* and *mixed* cycles. They consist of respectively only positive (as in J_1), only negative (as in J_3) and both kind of literals (as in J_2).

A justification is an argument for the truth value of its literals. Its *value* depends on the structure of the justification and the truth value of its leaves.

Definition 4 ($\mathcal{V}_I(J)$, the value of a justification). Let J be a justification, and I an interpretation.

- $\mathcal{V}_I(J) = \mathbf{f}$ if J contains either a leaf L with $I(L) = \mathbf{f}$ or a positive cycle.
- $\mathcal{V}_I(J) = \mathbf{u}$ if $\mathcal{V}_I(J) \neq \mathbf{f}$ and J contains either a leaf L with $I(L) = \mathbf{u}$ or a mixed cycle (or both).
- $\mathcal{V}_I(J) = \mathbf{t}$ otherwise (all leaves are \mathbf{t} and cycles, if any, are negative).

Example 5. Continuing Example 4. Independent of I , $\mathcal{V}_I(J_0) = \mathbf{t}$ and, because J_1 has a positive cycle, $\mathcal{V}_I(J_1) = \mathbf{f}$. Now, let $I = \{P^t, Q^t, R^f, S^t\}$. $\mathcal{V}_I(J_2) = \mathbf{u}$, indeed, although the only leaf (S) is \mathbf{t} , J_2 contains a mixed cycle; and finally $\mathcal{V}_I(J_3) = \mathbf{t}$, because the only leaf is \mathbf{t} and the only cycle is negative.

The value of a total justification J depends only on its cycles and the interpretation of its open symbols. J_L , the *restriction of J to L* , denotes the subgraph with L as root. A defined literal L is *justified* in a justification J if J_L is non-empty and total. The *supported value* $SV_I(L)$ of a defined literal L in an interpretation I is the maximal value in the truth order of its justifications i.e., $SV_I(L) = \max_{\leq} \{\mathcal{V}_I(J) \mid J \text{ is a justification and } L \text{ is justified in } J\}$.

It was proven in [9] that in the well-founded model M , the interpretation of defined literals agrees with the supported values, i.e., $M(L) = SV_M(L)$. Furthermore, [5] proved that $SV_M(L) = SV_M(\neg L)^{-1}$ for any defined literal L .

Example 6. Continuing from Example 4, let I' be an interpretation with $I'(S) = \mathbf{f}$. Then $SV_{I'}(P) = \mathbf{f}$, since for any justification J in which P is justified, either S is a leaf in J_P , or $J_P = J_1$. And indeed, let J' be the total justification obtained from J_3 by adding the edge $(\neg R, \neg S)$; then $SV_{I'}(\neg P) \geq \mathcal{V}_{I'}(J'_{\neg P}) = \mathcal{V}_{I'}(J') = \mathbf{t}$.

Hence one can compute the well-founded model by computing for each literal its supported value. However, searching over all justifications for the best one is infeasible. Fortunately, attention can be restricted to a subclass of justifications.

Definition 5 ((strict) support, positive residue). *Let J be a justification and I an interpretation. J supports I if $I(L) \leq_p I(\bigwedge Ch_J(L))$ and J strictly supports I if $I(L) = I(\bigwedge Ch_J(L))$ for each internal node L of J .*

The positive residue of J in I consists of the atoms that would be strictly supported if true but are undefined, i.e., $I(L) = \mathbf{u}$ and $I(\bigwedge Ch_J(L)) = \mathbf{t}$.

Definition 6 (Cycle-safe). *Let J be a justification and I an interpretation. J is cycle-safe in I if J contains neither mixed cycles with literals that are true in I nor positive cycles.*

Definition 7 (v-total). *Let J be a justification, I an interpretation, and \mathbf{v} a truth value, i.e., $\mathbf{v} \in \{\mathbf{f}, \mathbf{u}, \mathbf{t}\}$. Then J is \mathbf{v} -total in I if for each defined literal L such that $I(L) = \mathbf{v}$, L is justified in J .*

Note that if a justification J is \mathbf{f} -, \mathbf{u} - and \mathbf{t} -total w.r.t. an interpretation I , then J is total, but a total justification might not be \mathbf{f} -, \mathbf{u} - or \mathbf{t} -total w.r.t. any interpretation, because not every literal occurs in J .

Theorem 1. *Let I be an interpretation and J a justification. If the following conditions hold:*

- (i1) J strictly supports I ;
- (i2) J is cycle-safe in I ;
- (i3) J is \mathbf{t} -total in I ;
- (i4) J is \mathbf{u} -total in I ,

then for every defined literal L it holds that $I(L) = SV_I(L)$.

Proof. Let L be a defined literal. If $I(L) = \mathbf{t}$, then J_L is total. By (i1) we find that $\forall K \in Ch_J(L) : I(K) = \mathbf{t}$, and applying this argument on all defined literals in J_L , we find that $\forall M \in J_L : I(M) = \mathbf{t}$. Because of (i2), J_L does not contain positive or mixed cycles, so we find that $\mathcal{V}_I(J_L) = \mathbf{t}$, and therefore also that $SV_I(L) \geq \mathcal{V}_I(J_L) = \mathbf{t} = I(L)$.

If $I(L) = \mathbf{u}$, then also $I(\neg L) = \mathbf{u}$; both J_L and $J_{\neg L}$ are total. By (i1) we find that $\neg(\exists K \in Ch_J(L) : I(K) = \mathbf{f})$; again applying this argument inductively, we find that J_L contains no false literals. Since J_L also contains no positive cycles, $SV_I(L) \geq \mathcal{V}_I(J_L) \geq \mathbf{u}$. Symmetrically, also $SV_I(\neg L) \geq \mathbf{u}$ holds. Therefore, $SV_I(L) = SV_I(\neg L) = I(L) = \mathbf{u}$.

Finally, if $I(L) = \mathbf{f}$, then $I(\neg L) = \mathbf{t}$, so we find that $I(L) = I(\neg L)^{-1} = SV_I(\neg L)^{-1} = SV_I(L)$, and therefore $I(L) = SV_I(L)$.

Corollary 1. *Let I and J satisfy (i1)-(i4). Then $wfm_D(I|_{Open(D)}) = I$.*

In the next section we introduce an algorithm that incrementally constructs and maintains a 3-valued interpretation I and a justification J that satisfy the conditions (i1)-(i4).

4 Algorithm

4.1 Structure of the algorithm

The progenitor of our algorithm is the DLL algorithm [3], which is also the basis of most of today's SAT solvers. The DLL algorithm incrementally constructs an interpretation I that satisfies an input CNF theory T . In each stage, a choice literal is selected and is made true, after which *boolean propagation* is applied on the current assignment to infer other true literals. In particular, *unit propagation* makes the last non-false literal of a clause true. In case of conflict, backtracking returns to the last choice point. There, the alternative choice is made.³

Algorithm 1, the backbone of our algorithm is similar to the DLL algorithm. The input is a PC(ID) theory in MIDL normal form. However, the choice literal is an open one and boolean propagation is extended to include propagations according to the well-founded semantics by maintaining a 3-valued assignment I and a justification J . The aim of Steps 1 and 5 is to establish the conditions (i1)-(i4), as well as the following invariant:

- (i5) For each clause $\mathcal{C} \in T$, $I(\mathcal{C}) \geq \mathbf{u}$;

If so, obviously, I is a model when Step 11 returns SATISFIABLE. The ‘‘Boolean Propagation’’ (Step 5) is the core of the algorithm; it consists of two components:

Direct Propagation The input consists of I and J , the current interpretation and justification and the set of literals \mathcal{S} to be made true. The latter are made true, unit propagation on the clauses and propagation from body to head in the rules of the definitions are applied. This establishes the invariants (i1)-(i3) and (i5) unless a conflict is detected and backtracking is initiated.

³ UNSATISFIABLE is returned when no more backtrackings are possible.

```

1 Initialize  $I$  and  $J$  (see Section 4.4);
2 while there is an open atom  $A$  with  $I(A) = \mathbf{u}$  do
3   Select open atom  $A$  with  $I(A) = \mathbf{u}$ ;
4   Choose  $\mathcal{S} := \{A\}$  or  $\mathcal{S} := \{\neg A\}$ ;
   % Boolean Propagation:
5   while true do
6     if  $\mathcal{S} \neq \emptyset$  then
7        $\lfloor$  Direct_Propagation( $I, J, \mathcal{S}$ ); % Can initiate backtracking
8       if there is an  $L$  with  $I(L) = \mathbf{u}$  and  $Ch_J(L) = \emptyset$  then select such a  $L$ ;
9       else Break;
10       $\lfloor$   $\mathcal{S} := \text{Justify}(I, J, L)$ ;
11 if  $I$  is 2-valued then Return SATISFIABLE;
12 else Backtrack;

```

Algorithm 1: MIDL(T)

Justify The input consists of the current I and J , and a literal L that violates invariant (i4) (\mathbf{u} -totality) (L is a unknown defined literal that has no justification). The procedure tries to adjust J into a justification J' in which all literals of J'_L are strictly supported. In case this fails, an *unfounded set* [29] is found. The involved atoms must become false, their negation is returned in the set \mathcal{S} for processing by the direct propagation.

The \mathbf{u} -totality invariant is satisfied when the while loop exits in Step 9. It terminates because Direct_Propagation extends the interpretation I and Justify either reduces the number of unknown defined literals without justification or finds a unfounded set that leads to an extension of I in the next iteration. Hence:

Theorem 2. *Steps 5-10 of Algorithm 1 terminate; when they do via Step 9 then invariants (i1)-(i5) are satisfied.*

4.2 Direct Propagation (Algorithm 2)

The Direct_Propagation algorithm uses the following data structures:

- A data structure dj which associates with each defined literal L a direct justification of L . While $dj(L)$ is constant for conjunctive atoms and the negation of disjunctive atom, it consists of a selected literal (the *watched* literal) in the other cases.
- A boolean data structure $just$ over defined literals, indicating whether or not the literal's direct justification belongs to the current justification⁴.

The data structures dj and $just$ together determine a justification J in the following way: if $just(L)$, then $Ch_J(L) = dj(L)$; otherwise, $Ch_J(L) = \emptyset$.

Making a literal in \mathcal{S} true (Step 5) invalidates invariant (i1); however, the following weaker invariants are maintained:

⁴ Hence L is justified if $just(L)$ and the literals in $Ch_J(L)$ are justified

```

Assert: (i1), (i2), (i3) and (i5)
1 repeat
2   Pop  $L$  from  $\mathcal{S}$ ;
3   if  $I(L) = \mathbf{f}$  then backtrack (to choice point in Alg. 1);
4   if  $I(L) \neq \mathbf{t}$  then
5      $I(L) := \mathbf{t}$ ;
6     if  $L$  is defined then Push its original on  $\mathcal{S}$  (using the equivalences);
7     if  $L$  is open then for every clause  $c$  of  $T$  containing  $\neg L$  do
8       | if  $c$  has exactly one non-false literal  $L'$  then Push  $L'$  on  $\mathcal{S}$ ;
9     for every unknown literal  $L'$  for which  $\{L\}$  is a direct justification do
10    |  $just(L') := \mathbf{t}$ ;  $dj(L') := \{L\}$ ; Push  $L'$  on  $\mathcal{S}$ ;
11    for every unknown literal  $L'$  such that  $\neg L \in dj(L')$  do
12    | case  $L' = P$  with  $P \leftarrow \neg L \wedge \dots \in T$ , or  $L' = \neg P$  with
13    |    $P \leftarrow L \vee \dots \in T$ 
14    |   |  $dj(\neg L') := \{L\}$ ;  $just(\neg L) := \mathbf{t}$ ; Push  $\neg L'$  on  $\mathcal{S}$ ;
15    |   case  $L' = \neg P$  with  $P \leftarrow L \wedge \dots \in T$ 
16    |   | if  $P$  has no unknown body literals left then Push  $P$  on  $\mathcal{S}$ ;
17    |   | else
18    |   |   Select an unknown body literal  $L''$  of  $P$ ;  $dj(L') := \{\neg L''\}$ ;
19    |   |    $just(L') := \mathbf{t}$ ;
20    |   | case  $L' = P$  with  $P \leftarrow \neg L \vee \dots \in T$ 
21    |   | if  $P$  has no unknown body literals left then Push  $\neg P$  on  $\mathcal{S}$ ;
22    |   | else
23    |   |   if  $P$  has unknown open or negative body literal  $L''$  then
24    |   |   |  $dj(L') := \{L''\}$ ;  $just(L') := \mathbf{t}$ ;
25    |   |   | else
26    |   |   |   Select an unknown defined body atom  $Q$ ;
27    |   |   |    $dj(L') := \{Q\}$ ;  $just(L') := \mathbf{f}$ ;
28  until  $\mathcal{S} = \emptyset$ ;

```

Algorithm 2: Direct Propagation(I, J, \mathcal{S})

- (i6) J supports I ;
- (i7) the positive residue of J in I is a subset of \mathcal{S} ; for all defined literals $L \in \mathcal{S}$: $\mathcal{V}_{J_L}(I) = \mathbf{t}$; i.e., there exists a justification for making L true.

From invariant (i7), it follows that there is no positive residue when \mathcal{S} becomes empty, hence (i1) is restored. Also (i2) and (i3) are preserved:

t -totality (i3): A defined literal can only be made true in step Step 5 of Algorithm 2.⁵ Therefore, if a defined literal L has $I(L) = \mathbf{t}$, we know by (i6) that it has children in J that are all also true. By induction over all literals in J_L , J_L must be total and invariant (i3) holds.

⁵ This explains both the restriction in Step 3 of Algorithm 1 to *open* literals, and the requirement in Definition 1 that the CNF formula cannot contain defined literals.

Cycle-safeness (i2): By (i3) and (i6), true literals are justified, hence taking a true literal as the direct justification of a unknown literal cannot create a mixed cycle containing true literals. We must also verify that none of the steps in Algorithm 2 creates a positive cycle in J . Steps 2-10 don't change J . In the case of Step 12 $Ch_J(L')$ is only changed for a *negative* literal L' , in the case of Step 14 J is either not changed, or only for a negative literal. Finally, in the case of Step 18 a positive cycle might be created in J when the direct justification of the atom P is changed to a positive defined atom. This is prevented by removing it from J ($just(L') := \mathbf{f}$ implies $Ch_J(L') := \emptyset$). Consequently, \mathbf{u} -totality, invariant (i4), is violated.

Making an unknown literal true can violate invariant (i5) on the clauses. Step 8 not only performs unit propagation (not all clauses containing $\neg L$ are inspected but the Two Watched Literal technique is used) but also restoration of the invariant. By pushing the last non-false literal on the stack, a violation of the invariant will finally result in the uncovering of a conflict in Step 3.

Example 7. Let $T = \{\{P^D \leftarrow Q^D \vee A, Q^D \leftarrow P^D \wedge R^D \wedge \neg B, R^D \leftarrow C\}, P^D \equiv P, Q^D \equiv Q, R^D \equiv R\}$, $I = \emptyset$, and

$$J = \left[\begin{array}{ccccc} P^D \leftarrow Q^D \rightarrow R^D & \neg P^D \xrightarrow{\quad} \neg Q^D & \neg R & & \\ \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\ A & \neg B & C & \neg A & \neg C \end{array} \right]. \text{ Suppose Step 3 in Al-}$$

gorithm 1 chooses $\mathcal{S} := \{B\}$. Then Step 5 makes B true; subsequently Step 13 changes $Ch_J(\neg Q^D)$ from $\neg P^D$ to B , and pushes $\neg Q^D$ on \mathcal{S} . In a next iteration of the repeat loop, $\neg Q^D$ is made true, and $\neg Q$ is pushed on \mathcal{S} in Step 6.

Suppose that originally $\mathcal{S} := \{\neg A\}$ is chosen. After making A false, the algorithm ends up in Step 23. $dj(P^D)$ is changed from A to Q^D . In addition $just(P^D)$ is set to \mathbf{f} . Note that otherwise a positive cycle $P^D \leftrightarrow Q^D$ would have been created.

The Direct_Propagation algorithm can be understood as performing a watched literal technique for rules, similar to the Two Watched Literals (2WL) technique in SAT [20]. Every literal in a singleton direct justification (i.e., of a positive disjunctive atom, or of a negative conjunctive atom) is “watched”: when the literal becomes false, the corresponding rule has to be visited. When any other literal in that body becomes false, we don't need to visit the rule yet, since the “watched literal” is still unknown and thus might still justify the head.

Interpreting the head of each rule as a second watched literal, it can be seen that the Direct_Propagation algorithm actually has the same behaviour as the 2WL scheme has on the completion of the definition.

4.3 Justify (Algorithm 3)

Algorithm 3 tries to adjust J so that J_P is a total justification for the unknown atom P . If $dj(P)$, the direct justification of P , can be added to the current justification J (by setting $Just(P)$ to \mathbf{t}) without creating a positive cycle (involving

<p>Assert: (i1), (i2), (i3) and $Ch_J(P) = \emptyset$ Result: a set of negative literals \mathcal{S};</p> <pre> 1 if <i>IsCycleSafe</i>(P, dj) then $just(P) := t$; Return \emptyset; 2 $\mathcal{E} := \text{FindSet}(I, P)$; 3 $\mathcal{B} := \text{FindBottomSeeds}(I, \mathcal{E})$; $\mathcal{E} := \mathcal{E} \setminus \mathcal{B}$; 4 while $\mathcal{B} \neq \emptyset$ do 5 Select Q from \mathcal{B}; $\mathcal{B} := \mathcal{B} \setminus \{Q\}$; 6 for disjunctive atoms $Q' \in \mathcal{E}$ such that Q occurs in the body of Q' do 7 $dj(Q') := \{Q\}$; $just(Q') := t$; Move Q' from \mathcal{E} to \mathcal{B}; 8 for conjunctive atoms $Q' \in \mathcal{E}$ such that $dj(\neg Q') = \{\neg Q\}$ do 9 if there exists a $Q'' \in \mathcal{E}$ in the rule of Q' then $dj(\neg Q') := \{\neg Q''\}$; 10 else $just(Q') := t$; Move Q' from \mathcal{E} to \mathcal{B}; 11 Return $\{\neg Q \mid Q \in \mathcal{E}\}$; </pre>
--

Algorithm 3: Justify(I, J, P)

P) then we are done. This is tested in Step 1. This is a fairly simple test; however, if it fails, we will have to adjust the direct justification of P to construct a total justification.

The next step then is to find an overestimation of all unknown defined atoms that can potentially contribute to some J_P that is a total justification of P . Note that the unknown open literals and unknown negative literals can be used as valid leaves in a total justification, so they need not be included. This overestimation is computed in Step 2 and stored in \mathcal{E} . It consists of all atoms *reachable from P* in the dynamic positive dependency graph. The latter is defined as follows:

Definition 8 (dynamic positive dependency graph). *Let D be a definition in MIDL normal form, I a partial interpretation. The dynamic positive dependency graph of D in I is a directed graph (V, E) of atoms. $P \in V$ iff P is defined in D and $I(P) = \mathbf{u}$. $(P, Q) \in E$ iff $\{P, Q\} \subseteq V$ and P is the head of a rule and the atom Q occurs in its body.*

Observe that \mathcal{E} not only includes atoms without direct justification such as P itself, but also atoms with direct justification. In order to justify P , however, it may be necessary to update the direct justification of such atoms. In other words, \mathcal{E} is the set of *endangered atoms*: atoms, whose direct justification may have to be revised.

In the following steps, the algorithm tries to construct justifications for elements in \mathcal{E} . In Step 3, it collects in \mathcal{B} elements from \mathcal{E} that have a trivial justification:

- conjunctive atoms $C \in \mathcal{E}$ which have no elements from \mathcal{E} in their body (the unknown literals in their body are either open or negative)
- disjunctive atoms $D \in \mathcal{E}$ which have an unknown open or negative literal L in their body (their direct justification is set to that literal).

In the while loop, elements from \mathcal{B} are used one by one as seeds to construct justifications for other elements in \mathcal{E} . In Step 6, the direct justification of a disjunctive atom is set to the seed and in Step 8, a conjunctive atom is justified and removed from \mathcal{E} once it has no other body atoms in \mathcal{E} . In both steps, the newly justified atoms are added to the seeds. Observe that negative literals returned in Step 11 are justified; their direct justifications have been set in Step 9.

When exiting the while loop, all seeds have been used and what remains is an unfounded set of atoms. They have to be made false, hence their negation is returned by Justify.

This algorithm is an adaptation from [16] to the context of justifications.

Example 8. Continuing from Example 7, where we initially chose $\mathcal{S} := \{\neg A\}$. In Step 10 of Algorithm 1, $\text{Justify}(I, J, Q^D)$ will be called. Since $\text{dj}(Q^D) = \{P^D\}$ and $\text{dj}(P^D) = \{Q^D\}$, $\text{IsCycleSafe}(Q^D, \text{dj})$ fails. $\mathcal{E} = \{P^D, Q^D, R^D\}$ is computed in Step 2. In the next step we find $\mathcal{B} = \{R^D\}$ and therefore $\mathcal{E} = \{P^D, Q^D\}$, meaning that R^D might still justify P^D and/or Q^D becoming true. Since R^D only occurs in a conjunctive body (Q^D), but $\text{dj}(\neg Q^D) \neq \{\neg R^D\}$, R^D cannot justify anything, and the while-loop stops. Finally Justify returns $\{\neg P^D, \neg Q^D\}$ which are the negated literals from the unfounded set $\{P^D, Q^D\}$.

Let J be a justification, then we define $\mathcal{T}(J)$ as the set $\{L \mid L \in J \wedge \text{Ch}_J(L) = \emptyset \wedge L \text{ is defined}\}$, i.e. the defined leaves of J .

Proposition 1. *Let I , J and P satisfy the requirements in Algorithm 3. Let J' be the justification after termination of Algorithm 3 and \mathcal{S} the resulting set. Then either \mathcal{S} is empty and $\mathcal{T}(J') \subseteq \mathcal{T}(J) \setminus \{P\}$, or \mathcal{S} is not empty and contains the negated literals from an unfounded set: for each $\neg Q \in \mathcal{S}$, $\neg Q$ is justified and $\mathcal{V}_{J-\neg Q}(I) = \mathbf{t}$.*

Proposition 2. *Algorithm 3 preserves invariants (i1), (i2), (i3) and (i5).*

As shown in the following proof, these propositions formalize our claim in Section 4.1 that after each call to Justify either I is constant and the number of unknown defined leaves of J decreases, or I can be extended by the set returned by Justify.

Proof. of Theorem 2

It is given that interpretation I and justification J satisfy invariants (i1)-(i5). Also the following is given:

- Algorithm 2 and Algorithm 3 successfully terminate or returns conflict;
- after successful termination of Algorithm 2, invariants (i1), (i2), (i3) and (i5) are satisfied;
- Proposition 1 is satisfied;
- Proposition 2 is satisfied.

We have to show that the while loop of Algorithm 1 terminates (Steps 5-10), and that when it terminates via Step 9, invariants (i1)-(i5) are satisfied.

We prove termination by showing that any sequence of tuples $\langle I, J \rangle$ produced by the loop is decreasing in the following well-founded order \prec : $\langle I_1, J_1 \rangle \prec \langle I_2, J_2 \rangle$ iff either $I_2 \succ_p I_1$, where \succ_p is the precision order, or $I_1 = I_2$ and $\|\mathcal{T}(J_2)\| < \|\mathcal{T}(J_1)\|$.

After Step 7 in Algorithm 1, I has become strictly more precise, so $\langle I, J \rangle$ is strictly smaller w.r.t. \prec . After a call to “Justify” in Step 10, by Proposition 1 we know that there are two possibilities:

- \mathcal{S} is empty and $\mathcal{T}(J') \subseteq \mathcal{T}(J) \setminus \{P\}$: in this case clearly $\langle I, J \rangle$ has strictly decreased in the well-founded order; or
- \mathcal{S} is not empty, and for each $Q \in \mathcal{S}$, $\neg Q$ is justified and $I(J_{\neg Q}) = \mathbf{t}$: in this case, the next step in Algorithm 1 will be Direct Propagation, i.e., I will get strictly more precise *after* this step.

Therefore the while loop monotonically decreases $\langle I, J \rangle$ w.r.t. \prec . Now, any sequence $\langle I_1, J_1 \rangle \succ \langle I_2, J_2 \rangle \succ \dots$ has a lower bound $\langle I_u, J_u \rangle$: namely I_u is two-valued and $I_u \geq_p I_i$ for any i , and $\mathcal{T}(J_u)$ empty. This proves termination of the loop.

Both Algorithm 2 and 3 satisfy invariants (i1), (i2), (i3) and (i5), so clearly these are also satisfied after successful termination of the while loop.

Remains to prove (i4) (**u**-totality). Successful termination only occurs when there is no defined literal L with $I(L) = \mathbf{u}$ and $Ch_J(L) = \emptyset$. Take a defined literal L with $I(L) = \mathbf{u}$; then by invariant (i1), $\neg(\exists K \in Ch_J(L) : I(K) = \mathbf{f})$. Applying this argument inductively on all defined literals in J_L , we find that J_L does not contain any false literals. Therefore, J_L is total.

4.4 Correctness, initialization and optimization issues

Soundness and Completeness. Soundness has been argued in Section 4.1: I is a model when Algorithm 1 returns SATISFIABLE. Completeness follows from the completeness of the DLL algorithm and the observation that our Boolean Propagation extends unit propagation of DLL and computes after each choice point the well-founded model extending the interpretation of the open literals.

Initialization. Step 1 of Algorithm 1 should initialize an interpretation I and a partial justification J such that invariants (i1)-(i5) are satisfied. To do this, we initialize I to the empty interpretation and $dj(L)$ to an arbitrary direct justification of L , for every defined literal. To avoid positive cycles, we set $just(P) := \mathbf{f}$ for every defined atom P . Then the set of unit clauses is collected in \mathcal{S} , and the initialisation executes the same while loop as described in Step 5 of Algorithm 1.

Backtracking. Backtracking restores I and J to some old value, say I_b and J_b . The invariants (i1)-(i5) hold for I_b and J_b . In fact, with exception of (i4), they also hold for I_b and J . Indeed, a conflict arises at a moment when (i2), (i3) and (i6) all hold. After restoring I to a previous level, for every defined literal

L it holds that $I(L) = I(\bigwedge dj(L))$, and hence (i6) still holds, i.e., J supports I . Since (i1)-(i3), (i5) are sufficient pre-conditions for Boolean Propagation, the justification need not be restored upon backtracking.

Endangered Atoms. After an unsuccessful cycle-safety check (which is in fact optional), the first step of Justify is a call to FindSet to produce \mathcal{E} , the set of atoms endangered by the modified direct justification of P . There are a number of alternative possibilities. For example, the algorithm would remain correct if FindSet would return all unknown atoms in $Def(D)$ instead. However, one could also compute a smaller set of endangered atoms: Those in the strongly connected component (SCC) of P in the dynamic positive dependency graph. An SCC is a set \mathcal{S} of atoms, such that every element of \mathcal{S} can reach every other element of \mathcal{S} via the directed edges. Tarjan’s algorithm searches for strongly connected components in time linear in the size of the graph [26, 22]. Also the SCC of the static positive dependency graph, which we define as the dynamic positive dependency graph after the initialization (Step 1 in Algorithm 1), could be used. The smaller the set of endangered atoms, the faster it is processed. However, there is a trade-off between the work spent to compute the endangered atoms and that spent on processing them.

5 Experimental results

We have made a prototype implementation, called MIDL, of the algorithm described in the previous section. We have compared MIDL to Smodels [21] and idsat(zChaff) [23, 30]. We have thus a representative for each of the three approaches mentioned in Section 1. Note that the input to Smodels is hand-coded and hence behaves possibly better than what one would obtain by an automatic mapping from PC(ID) to ASP. We used two classes of problems: N -queens and Hamiltonian cycles⁶. All the experiments were run on 863 MHz P-III with 254 MB of RAM. In Table 1 timings to find the first model (averaged over 5 runs) are given in seconds.

Since MIDL doesn’t enjoy the fine-tuned optimizations of a state of the art SAT or ASP solver, we don’t expect comparable efficiencies. Most notably, MIDL’s data structures are not yet minimized, which can be expected to lead to serious losses in cache usage. Also its heuristics are very crude.

Indeed, the results show poor scaling. However, observe that MIDL outperforms the other solvers on some of the smaller problems. Remarkably, zChaff cannot cope with several Hamiltonian cycle problems; as they contain a lot of induction, idsat increased the size of these problems 20-fold. This confirms the usefulness of our native approach with respect to a translation to SAT. For an important class of problems, MIDL is currently the best PC(ID) model generator.

⁶ We denote the Hamiltonian cycle problems by “H-#vertices-#nodes”. Encodings are taken from <http://asparagus.cs.uni-potsdam.de/>; randomly generated graphs were used.

	MiDL	idsat+zCh	Smod.
9-queens	0.08	0.56+0.02	0.05
11-queens	1.49	0.79+0.21	0.18
13-queens	10.77	1.18+0.13	0.43
15-queens	296.91	1.67+0.16	1.85
H-20-200	0.32	61.19+7.62	0.10
H-25-200	0.06	96.5+82.2	0.10
H-30-200	1.97	98.5+137	0.13
H-35-200	3.65	#+#	0.18
H-20-400	0.12	72.87+#	0.35
H-25-400	0.20	128.0+#	0.55
H-30-400	0.18	209.2+#	0.56
H-35-400	#	#+#	0.87

Table 1. Timings (*sec*) of MiDL, idsat using zChaff, and Smodels. # = >10min.

		MiDL	MiDL-	MiDL	MiDL
			SCC	SCC-	SCC-
H-20-200	<i>U</i>	1427	1461	1022	1433
	<i>E</i>	4570	20327	2535	11363
	time	0.32	0.34	3.52	1.57
H-30-200	<i>U</i>	1933	2013	474	1976
	<i>E</i>	10567	116186	2708	62460
	time	1.97	2.87	1.20	3.62
H-20-400	<i>U</i>	248	283	112	179
	<i>E</i>	1112	6573	785	4435
	time	0.12	0.17	0.25	0.24
H-30-400	<i>U</i>	209	231	85	180
	<i>E</i>	2002	31618	2031	22975
	time	0.18	1.70	0.32	1.77

Table 2. Comparison of unfounded set sizes (*U*) vs. number of “endangered atoms” (*E*) and timings (*sec*) for different variants of MiDL

In Table 2 we compare timings⁷, the total number of atoms found to be in an unfounded set (*U*) and the total number of endangered atoms considered (*E*) for some variants of MiDL: in MiDL- and MiDL SCC-, the optional Step 1 of Algorithm 3 is disabled, in MiDL SCC and MiDL SCC- the FindSet procedure returns the SCC of the dynamic positive dependency graph, as described in Section 4.4. The first important observation is that the cycle-safety check almost consistently yields faster results. A second observation is that MiDL SCC is doing more intelligent work: it uses less endangered atoms, and it also finds much less unfounded sets, i.e., finds falsity of atoms through Direct Propagation more often. Still, MiDL outperforms MiDL SCC, suggesting that a more careful implementation that removes some overhead of the SCC computation might be beneficial.

6 Conclusions, related and future work

This work is one of the first attempts to build a SAT(PC(ID)) system. We have chosen for a direct implementation, in contrast to a mapping to ASP, or to propositional logic, as was done in [23]. The latter approach is similar to ASP systems such as ASSAT [14] and Cmodels [13].

Despite semantical differences between PC(ID) and ASP, the algorithms presented here share a lot of structure with those of Smodels and DLV. The main novelties of our approach come through the use of justifications:

- this enables us to do the beneficial cycle-safety check;
- it integrates nicely with a watched literal technique for rules;

⁷ To save space, we’ve selected but a few example problems; other problems exhibit the same behaviour.

- the justification graph can be seen as a straightforward extension of the implication graph, which is used for Clause Learning [18]. In the near future, we intend to include this important SAT technique in MIDL.

Another strongly related system is `dcs` [11]. This system can be viewed as a model generator for a fragment of FO(ID). The system takes as input a function free FO theory and an inductive definition consisting of Horn rules, and computes Herbrand models of this theory.

In future work, we plan to re-implement MIDL, investigating a variety of techniques from SAT, logic programming and ASP. Interesting recent optimizations to SAT solvers are described in [20, 24, 18, 30]. Potentially relevant techniques for computing the well-founded semantics are described in [28, 1, 16, 25].

Finally, we mention some issues that will be investigated in the near future:

- The current solver uses rules only in a bottom up propagation. In some situations it is definitely worthwhile to also exploit propagation from head to body.
- As shown in [20, 24], the quality of the search algorithm strongly depends on the heuristics. Now that our search algorithm is more or less fixed, we should start to evaluate different heuristics for the system.
- A major task in the project is to build an efficient grounder which reduces a FO(ID) theory to a propositional theory by grounding it with respect to the Herbrand universe.

References

1. K.A. Berman, J.S. Schlipf, and J.V. Franco. Computing the well-founded semantics faster. In V. Marek and A. Nerode, editors, *LPNMR*, volume 928 of *Lecture Notes in Computer Science*, pages 113–126. Springer, 1995.
2. R.J. Brachman and H.J. Levesque. Competence in knowledge representation. In *Proc. of the National Conference on Artificial Intelligence*, pages 189–192, 1982.
3. M. Davis, G. Longemann, and D. Loveland. A machine program for theorem proving. *Communications of the ACM*, 5:394–397, 1962.
4. T. Dell’Armi, W. Faber, G. Ielpa, C. Koch, N. Leone, S. Perri, and G. Pfeifer. System description: DLV. In T. Eiter, W. Faber, and M. Truszczyński, editors, *LPNMR*, volume 2173 of *LNCS*, pages 424–428. Springer, 2001.
5. M. Denecker. *Knowledge Representation and Reasoning in Incomplete Logic Programming*. PhD thesis, Department of Computer Science, K.U.Leuven, 1993.
6. M. Denecker. The well-founded semantics is the principle of inductive definition. In J. Dix, L. Fariñas del Cerro, and U. Furbach, editors, *Logics in Artificial Intelligence*, volume 1489 of *LNAI*, pages 1–16. Springer-Verlag, 1998.
7. M. Denecker. Extending classical logic with inductive definitions. In J. Lloyd et al., editor, *CL*, volume 1861 of *LNAI*, pages 703–717. Springer, 2000.
8. M. Denecker, M. Bruynooghe, and V. Marek. Logic programming revisited: logic programs as inductive definitions. *ACM Transactions on Computational Logic*, 2(4):623–654, 2001.
9. M. Denecker and D. De Schreye. Justification semantics: a unifying framework for the semantics of logic programs. In *LPNMR*, pages 365–379. MIT Press, 1993.

10. M. Denecker and E. Ternovska. Inductive situation calculus. In D. Dubois, C.A. Welty, and M. Williams, editors, *KR*, pages 545–553. AAAI Press, 2004.
11. D. East and M. Truszczyński. *dcs*: An implementation of datalog with constraints. *CoRR*, cs.AI/0003061, 2000.
12. M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–387, 1991.
13. Y. Lierler and M. Maratea. Cmodels-2: SAT-based answer set solver enhanced to non-tight programs. In V. Lifschitz and I. Niemelä, editors, *LPNMR*, volume 2923 of *LNCS*, pages 346–350. Springer, 2004.
14. F. Lin and Y. Zhao. ASSAT: computing answer sets of a logic program by sat solvers. *Artif. Intell.*, 157(1-2):115–137, 2004.
15. J.W. Lloyd and R.W. Topor. Making prolog more expressive. *Journal of Logic Programming*, 1(3):225–240, 1984.
16. Z. Lonc and M. Truszczyński. On the problem of computing the well-founded semantics. In J.W. Lloyd, V. Dahl, U. Furbach, M. Kerber, K. Lau, C. Palamidessi, L.M. Pereira, Y. Sagiv, and P.J. Stuckey, editors, *Computational Logic*, volume 1861 of *LNCS*, pages 673–687. Springer, 2000.
17. M. Mariën, D. Gilis, and M. Denecker. On the relation between ID-logic and answer set programming. In J.J. Alferes and J.A. Leite, editors, *JELIA*, volume 3229 of *LNCS*, pages 108–120. Springer, 2004.
18. J.P. Marques-Silva and K.A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Trans. Computers*, 48(5):506–521, 1999.
19. D.G. Mitchell and E. Ternovska. A framework for representing and solving NP search problems. In AAAI Press/MIT Press, editor, *Twentieth National Conf. on Artificial Intelligence (AAAI-05)*, 2005.
20. M.W. Moskewicz, C.F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *DAC*, pages 530–535. ACM, 2001.
21. I. Niemelä, P. Simons, and T. Syrjänen. Smodels: a system for answer set programming. In *NMR*, 2000.
22. E. Nuutila and E. Soisalon-Soininen. On finding the strongly connected components in a directed graph. *Inf. Process. Lett.*, 49(1):9–14, 1994.
23. N. Pelov and E. Ternovska. Reducing inductive definitions to propositional satisfiability. In *ICLP*, LNCS, 2005.
24. L. Ryan. Efficient algorithms for clause-learning SAT solvers. Master’s thesis, Simon Fraser University, 2004.
25. K.F. Sagonas, T. Swift, and D.S. Warren. XSB as an efficient deductive database engine. In R.T. Snodgrass and M. Winslett, editors, *SIGMOD Conference*, pages 442–453. ACM Press, 1994.
26. R.E. Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.
27. G.S. Tseitin. On the complexity of derivation in propositional calculus. *Automation of Reasoning (Classical Papers on Computational Logic, 1967-1970)*, 2:466–483, 1983.
28. A. Van Gelder. The alternating fixpoint of logic programs with negation. *Journal of Computer and System Sciences*, 47(1):185–221, 1993.
29. A. Van Gelder, K.A. Ross, and J.S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 38(3):620–650, 1991.
30. L. Zhang, C.F. Madigan, M.W. Moskewicz, and S. Malik. Efficient conflict driven learning in a boolean satisfiability solver. In *ICCAD*, pages 279–285. ACM, 2001.