TRANSFORMATION-BASED BOTTOM-UP COMPUTATION OF THE WELL-FOUNDED MODEL

STEFAN BRASS, JÜRGEN DIX, BURKHARD FREITAG AND ULRICH ZUKOWSKI

▷ We present a bottom-up algorithm for the computation of the well-founded model of non-disjunctive logic programs. Our method is based on the notion of conditional facts and elementary program transformations studied by Brass and Dix [BD94, BD97, BD98b]. However, their "residual program" can grow to exponential size, whereas for function-free programs our "program remainder" is always polynomial in the size of the extensional database (EDB). Our approach is also closely related to the alternating fixpoint procedure [VG89, VG93]. However, the alternating fixpoint procedure is known to have inefficiencies due to redundant recomputations of possible facts. By using conditional facts that can be deleted directly instead of recomputing the complement our approach is guaranteed to need not more work than the alternating fixpoint procedure and to be much more efficient in many cases.

The magic set transformation as a method for bottom-up query answering is known to have problems with undefined magic facts in the context of the well-founded semantics. However, our approach treats magic set transformed programs correctly, i.e. it always computes a relevant part of the well-founded model of the original program. We show that our approach is guaranteed to need not more work than the well-founded magic sets method [KSS95] or the magic alternating fixpoint procedure [Mor96] and is much more efficient for many programs. Again the use of conditional facts gives valuable information to select the right magic facts to use and to avoid redundant recomputations.

Address correspondence to Stefan Brass, School of Information Sciences, University of Pittsburgh, 135 North Bellefield Ave., Pittsburgh, PA 15260, USA, e-mail: sbrass@sis.pitt.edu and Jürgen Dix, University of Koblenz, Dept. of Computer Science, Rheinau 1, D-56075 Koblenz, Germany, e-mail: dix@mailhost.uni-koblenz.de and Burkhard Freitag, Ulrich Zukowski, Universität Passau, Fakultät für Mathematik und Informatik, D-94030 Passau, Germany, e-mail: {freitag,zukowski}@fmi.uni-passau.de

THE JOURNAL OF LOGIC PROGRAMMING

© Elsevier Science Inc., 1994
655 Avenue of the Americas, New York, NY 10010 0743-1066/94/$7.00

These results show that our approach is a valuable tool to analyze, compare, and optimize existing evaluation methods or to create new strategies that are automatically proven to be correct if they can be described by a sequence of transformations of our framework. We have developed a prototypical implementation allowing experiments that show that the theoretical results carry over to practical applications.

Keywords: well-founded semantics, bottom-up evaluation, alternating fixpoint, magic set transformation

◁

1. Introduction

The next generation of deductive database systems (see RAMRISHNAN's article in [AMTW99]) will probably support the full class of normal programs. It seems also very likely that the well-founded semantics will be chosen by nearly all system designers, because it has a unique model. Furthermore, the time complexity of computing the well-founded model of a given intensional database (IDB) is polynomial in the size of the extensional database (EDB) [VGRS91]. In contrast, it has been shown by MAREK and TRUSZCZYNSKI [MT91] that even for a propositional logic program P, determining whether P has a stable model is NP-complete. Extensions of deductive database systems that can deal with the well-founded semantics are already realized in XSB [CW96, SSW94, CSW95] and announced for LOLA [FSS91, ZF96a, ZF96b, ZF97].

The SLG-resolution of CHEN and WARREN [CW93, CW96, CSW95], as implemented in the XSB system, is the most prominent *top-down* method for the computation of the well-founded model of a normal program. In this paper we characterize "good" *bottom-up* methods in terms of elementary program transformations. Essentially, the bottom-up algorithms that compute the well-founded model of a normal program are:

- The alternating fixpoint approach, introduced by VAN GELDER [VG89, VG93] and further developed by KEMP, STUCKEY and SRIVASTAVA [KSS95].
- The computation of the residual program, suggested by BRY [Bry89, Bry90] and independently by DUNG/KANCHANASUT [DK89a, DK89b], and extended by BRASS and DIX [BD97, BD98c].

The alternating fixpoint procedure is known to have efficiency problems, since in every iteration many facts have to be recomputed. The residual program approach avoids recomputations of this kind. But it is possible that the residual program can grow to exponential size while for function-free programs the alternating fixpoint approach guarantees a number of derived facts that is polynomial in the size, i.e. the number of tuples, of the extensional database (EDB). However, the residual program contains important information which is not provided by the alternating fixpoint method.

It is natural to ask for an algorithm which combines the advantages of both bottom-up approaches. In this paper, we present a framework for the computation of the well-founded model that is based on the residual program method but guarantees polynomial complexity like the alternating fixpoint procedure. The residual

program method is based on elementary program transformations and the concept of *conditional facts*. Conditional facts are ground rules having exactly those negative literals in their bodies that can not (yet) be resolved because their complement is not (yet) known to be true or false. This can be regarded as a non-procedural equivalent of the *delay* operation needed in SLG-resolution [CW93, CW96, CSW95]. Like in SLG-Resolution, the key idea to avoid the exponential blow-up of the residual program is to delay not only negative literals but also positive literals which depend on delayed negative literals. The transformation-based *program remainder* method which we propose in this paper can be seen as a bottom-up equivalent of SLG-resolution for the case of range-restricted function-free programs.

But our method is not applicable only to range-restricted programs. In many cases, the magic set transformed version of a non-range-restricted program w.r.t. a given query is range-restricted. Although it is known that the magic set transformation has problems in the context of the well-founded semantics our method treats magic set transformed programs correctly. Moreover, we show that our approach is guaranteed to need not more work than the well-founded magic sets method [KSS95] or the magic alternating fixpoint [Mor96] and is much more efficient for many programs.

The rest of the paper is organized as follows. After introducing preliminaries in Section 2, we recall the alternating fixpoint procedure in Section 3. In Section 4, we introduce our program transportation approach for ground programs. In Section 5 we present an method that performs an intelligent grounding. Section 6 compares our approach to the alternating fixpoint procedure. Section 7 introduces a regular expression like syntax to express and compare different evaluation strategies. In Section 8 we extend our transformation approach to handle magic set transformed programs. Section 10 concludes the paper.

2. Preliminaries

A rule is of the form $A \leftarrow L_1 \wedge \cdots \wedge L_n$, where the head A is an atom and each body literal L_i is a positive literal, i.e. an atom, or a negative literal $L_i = \mathbf{not}\, B$. We treat the rule body as a set of literals and write also $A \leftarrow \mathcal{B}$ with $\mathcal{B} = \{L_1, \ldots, L_n\}$. Consequently, a fact A is represented by the rule $A \leftarrow \emptyset$. We also use the notation $A \leftarrow \mathcal{B} \wedge \mathbf{not}\, \mathcal{C}$ with atom sets $\mathcal{B}$ and $\mathcal{C}$.

A program is a set of rules as introduced above. We consider normal, i.e., non-disjunctive, logic programs without function symbols. We assume that all rules are range-restricted, i.e. that each variable of the rule appears also in a positive body literal. The *size* of a program is the number of literal occurrences.

Let P be a program. $BASE(P)$ denotes the Herbrand base of P, i.e. the set of all ground atoms. It might happen that our transformations eliminate some constants and predicates from the program, but we use the Herbrand base $BASE(P)$ of the original program. We write $ground(P)$ for the Herbrand instantiation of a program P. For a ground program P we define the following sets:

$$\begin{aligned} facts(P) &:= \{A \in BASE(P) \mid (A \leftarrow \emptyset) \in P\}, \\ heads(P) &:= \{A \in BASE(P) \mid \text{there is a } \mathcal{B} \text{ such that } (A \leftarrow \mathcal{B}) \in P\}. \end{aligned}$$

The complement of a literal L is denoted by $\sim L$, i.e., $\sim(B) = \mathbf{not}\, B$ and, conversely, $\sim(\mathbf{not}\, B) = B$. For a set S of literals, $\sim S$ denotes the set of the complements of

the literals in S. S is *consistent* if and only if $S \cap {\sim}S = \emptyset$. For a set S of literals we define the following sets:

$$pos(S) := \{A \in S \mid A \text{ is a positive literal } \},$$
$$neg(S) := \{A \mid \mathbf{not}\, A \in S\}.$$

Let P be a logic program. A *partial interpretation* for P is a consistent set I of ground literals such that

$$pos(I) \cup neg(I) \subseteq BASE(P),$$

i.e., its set of atoms is a subset of the Herbrand base of P. A *total interpretation* for P is a partial interpretation I such that

$$A \in BASE(P) \implies A \in I \text{ or } \mathbf{not}\, A \in I,$$

i.e., for each atom A of the Herbrand base of P either A or its complement is contained in I. Two partial interpretations are said to agree on an atom Q if they contain the same ground instances of Q and $\mathbf{not}\, Q$.

For an operator T we define $T \uparrow 0 := \emptyset$ and $T \uparrow i := T(T \uparrow i-1)$, for $i > 0$. $\mathrm{lfp}(T)$ denotes the least fixpoint of T, i.e. the smallest set S such that $T(S) = S$.

We rely on the definition of the *well-founded partial model* W_P^* of P as given in [VGRS91]. An operational characterization of W_P^* will be given in Section 3.

3. The Alternating Fixpoint Procedure

Let us recall the definition of the alternating fixpoint procedure. We introduce an extended version of the immediate consequence operator that uses two different sets of facts for positive and negative subgoals, respectively. Actually, this is an adaption of the stability transformation found in [VG89, VG93] to our purposes. This operator has been introduced and investigated by Przymusinski in [Prz89, Prz90, Prz91].

Definition 3.1. (Extended Immediate Consequence Operator) Let P be a normal logic program. Let I and J be sets of ground atoms. The set $T_{P,J}(I)$ of *immediate consequences of I w.r.t. P and J* is defined as follows:

$$T_{P,J}(I) := \{A \mid \text{there is } A \leftarrow \mathcal{B} \in ground(P) \text{ with } pos(\mathcal{B}) \subseteq I \text{ and } neg(\mathcal{B}) \cap J = \emptyset\}.$$

If P is definite, the set J is not needed and we obtain the standard immediate consequence operator T_P by $T_P(I) = T_{P,\emptyset}(I)$.

$T_{P,J}$ checks negative subgoals against the set of possibly true atoms that is supplied as the argument J. This allows the following elegant formulation of the alternating fixpoint procedure.

Definition 3.2. (Alternating Fixpoint Procedure) Let P be a normal logic program. Let P^+ denote the subprogram consisting of the definite rules of P. Then the sequence $(K_i, U_i)_{i \geq 0}$ with sets K_i of true (known) facts and U_i of possible (unknown) facts is defined by:

$$K_0 := \mathrm{lfp}(T_{P^+}) \qquad U_0 := \mathrm{lfp}(T_{P,K_0})$$
$$i > 0 : K_i := \mathrm{lfp}(T_{P,U_{i-1}}), \; U_i := \mathrm{lfp}(T_{P,K_i}).$$

The computation terminates when the sequence becomes stationary, i.e., when a fixpoint is reached in the sense that

$(K_j, U_j) = (K_{j+1}, U_{j+1})$.

This computation schema is called the *Alternating Fixpoint Procedure* (AFP).

Theorem 3.3 ([VG93]) Let the sequence $(K_i, U_i)_{i \geq 0}$ be defined as above. Then there is a $j \geq 0$ such that $(K_j, U_j) = (K_{j+1}, U_{j+1})$. The well-founded model W_P^* of P can be directly derived from the fixpoint (K_j, U_j), i.e.,

$$W_P^* = \{L \mid L \text{ is a positive ground literal and } L \in K_j \text{ or}$$
$$L \text{ is a negative ground literal } \mathbf{not}\, A \text{ and } A \in BASE(P) - U_j\}.$$

The following results which follow easily from the anti-monotonicity of $T_{P,J}$ w.r.t. J are also stated in [KSS95].

Lemma 3.4 Let the sequence $(K_i, U_i)_{i \geq 0}$ be defined as above. Then the following holds for $i \geq 0$: $K_i \subseteq K_{i+1}$, $U_i \supseteq U_{i+1}$, $K_i \subseteq U_i$.

Whereas the sets K_i of true facts can be computed incrementally, the sets U_i of possible facts have to be recomputed in every iteration step. It is well-known that this leads to many unnecessary recomputations even for simple programs.

Example 3.5. Consider the following logic program taken from [Mor96, KSS95]

$p(X) \leftarrow t(X, Y, Z), \mathbf{not}\, p(Y), \mathbf{not}\, p(Z)$.
$p(X) \leftarrow p_0(X)$.

and the following base facts:

$$p_0(c_2), t(a, a, b), t(b_1, c_1, b_2), t(b_2, c_2, b_3), \ldots, t(b_n, c_n, b_{n+1}).$$

Due to the linear character of the graph of dependencies between the atoms $p(b_1)$ to $p(b_n)$ we would expect that it is possible to compute the well-founded model of this program in linear time w.r.t. n. However, the alternating fixpoint procedure needs $\frac{n}{2}$ iterations, each deriving a number of facts that is linear in n. So the total time complexity is at least quadratic in n.

4. Basic Ground Transformations

Brass and Dix [BD94, BD97, BD98b, BD98c] have introduced a framework for studying and computing negation semantics by means of elementary program transformations. Based on this framework we introduce our transformation approach.

Definition 4.1. (Program Transformation) A *program transformation* is a relation $\mapsto$ between ground logic programs.

The following two transformations reduce negative body literals if their truth value is obvious. These transformation are also used in [BD98c].

Definition 4.2. (Positive Reduction) Let P_1 and P_2 be ground programs. P_2 results

from P_1 by *positive reduction* ($P_1 \mapsto_P P_2$) iff there is a rule $A \leftarrow \mathcal{B}$ in P_1 and a negative literal $\mathbf{not}\, B \in \mathcal{B}$ such that there is no rule about B in P_1, i.e., $B \notin heads(P_1)$, and $P_2 = \big(P_1 - \{A \leftarrow \mathcal{B}\}\big) \cup \big\{A \leftarrow (\mathcal{B} - \{\mathbf{not}\, B\})\big\}$.

Definition 4.3. (Negative Reduction) Let P_1 and P_2 be ground programs. P_2 results from P_1 by *negative reduction* ($P_1 \mapsto_N P_2$) iff there is a rule $A \leftarrow \mathcal{B}$ in P_1 and a negative literal $\mathbf{not}\, B \in \mathcal{B}$ such that B appears as a fact in P_1, i.e., $B \in facts(P_1)$, and $P_2 = P_1 - \{A \leftarrow \mathcal{B}\}$.

In [BD98c] positive body literals are resolved by unfolding. However, it is well-known that by unrestricted unfolding the resulting program can grow to a size that is exponential in the size of the original program [CW96, ZBF97]. To achieve a polynomial complexity we introduce the following two transformations that are special cases of unfolding and reduce positive body literals only if their truth value is obvious.

Definition 4.4. (Success) Let P_1 and P_2 be ground programs. P_2 results from P_1 by *success* ($P_1 \mapsto_S P_2$) iff there is a rule $A \leftarrow \mathcal{B}$ in P_1 and a positive literal $B \in \mathcal{B}$ such that $B \in facts(P_1)$, and $P_2 = \big(P_1 - \{A \leftarrow \mathcal{B}\}\big) \cup \big\{A \leftarrow (\mathcal{B} - \{B\})\big\}$.

Definition 4.5. (Failure) Let P_1 and P_2 be ground programs. P_2 results from P_1 by *failure* ($P_1 \mapsto_F P_2$) iff there is a rule $A \leftarrow \mathcal{B}$ in P_1 and a positive literal $B \in \mathcal{B}$ such that there is no rule about B in P_1, i.e., $B \notin heads(P_1)$, and $P_2 = P_1 - \{A \leftarrow \mathcal{B}\}$.

Definition 4.6. Let $\mapsto_{PSNF}$ be the rewriting system consisting of the above four transformations, i.e. $\mapsto_{PSNF} := \mapsto_P \cup \mapsto_N \cup \mapsto_S \cup \mapsto_F$.

This rewriting system is to be applied to a program as long as it is applicable, i.e., until a normal form w.r.t. these transformations is reached.

Definition 4.7. (Normal form) A program P' is a *normal form* of a program P w.r.t. a transformation $\mapsto$ iff $P \mapsto^* P'$, and P' is irreducible, i.e. there is no program P'' with $P' \mapsto P''$.

When a normal form P' is reached, the result of the transformation is interpreted in the following way: atoms appearing as facts in P' are true in P', atoms from the Herbrand base of P not appearing in any rule head of P' are false in P', all remaining atoms are undefined in P'.

Definition 4.8. (Known Literals) Let P be a ground program and let S be a set of ground atoms. The set of positive and negative ground literals with atoms in S having an obvious truth value in P is denoted by $known_S(P)$:

$$known_S(P) := \{L \mid L \in S \text{ is a positive ground literal and } L \in S \cap facts(P) \text{ or } L \text{ is a negative ground literal } \mathbf{not}\, A \text{ and } A \in S - heads(P)\}.$$

The next theorem states that the four transformations *success*, *failure*, *positive* and *negative reduction* correspond exactly to the Fitting operator Φ_P, which is the procedural counterpart of the 3-valued completion $comp_3$ (we refer to [Fit85, Dix95] for a detailed exposition of Φ_P).

Theorem 4.9 ($\mapsto_{PSNF}$ corresponds to $\mathrm{lfp}(\Phi_P)$)
The rewriting system $\mapsto_{PSNF}$ is confluent and strongly terminating. It therefore induces a normal form $fitt(P)$ for any program P. The known literals with respect to this normal form constitute exactly the least fixpoint $\mathrm{lfp}(\Phi_P)$ of Fitting's operator Φ_P:

$$\mathrm{lfp}(\Phi_P) = known_{BASE(P)}\big(fitt(P)\big).$$

PROOF. See Appendix A. □

Note, that for a given ground program P this normal form $fitt(P)$ can be computed in time that is linear in the size of P. An efficient algorithm based on [DG84] can be found in [NS96]. The four parts of this implementation, that reduce body literals, correspond directly to the four basic transformations defined above. Thus, that algorithm can be seen as an efficient implementation of our rewriting system $\mapsto_{PSNF}$. For many programs, e.g. for programs without positive loops, the fitting operator computes the well-founded model. This positive effect of the fitting operator has also been pointed out in [BSF95, SNV95]. However, positive loops have to be detected and removed separately.

Example 4.10. Consider the following program *Loop*:

$p.$
$q \leftarrow \mathbf{not}\, p.$
$q \leftarrow r.$
$r \leftarrow q.$

We can apply *negative reduction* to delete $q \leftarrow \mathbf{not}\, p$, since $\mathbf{not}\, p$ is obviously false. But $\mapsto_{PSNF}$ does not allow to delete $q \leftarrow r$ and $r \leftarrow q$.

To remove positive loops we introduce the following transformation.

Definition 4.11. (Loop Detection) Let P_1 and P_2 be ground programs. P_2 results from P_1 by *loop detection* ($P_1 \mapsto_L P_2$) iff there is a set $\mathcal{A}$ of ground atoms such that

1. for each rule $A \leftarrow \mathcal{B}$ in P_1, if $A \in \mathcal{A}$, then $\mathcal{B} \cap \mathcal{A} \neq \emptyset$,
2. $P_2 := \{A \leftarrow \mathcal{B} \in P_1 \mid A \notin \mathcal{A}\}$,
3. $P_1 \neq P_2$.

In contrast to the four transformations of $\mapsto_{PSNF}$ that can be applied locally, for *loop detection* we need a more global perspective on the program. However, the set $\mathcal{A}$ is the *greatest unfounded set* of P_1 (w.r.t. the empty interpretation $\emptyset$, cf. [VGRS91]). This greatest unfounded set consists of all positive ground atoms which are not possibly true, i.e., cannot be derived by assuming all negative literals to be true. According to Definition 3.1, the extended immediate consequence operator $T_{P_1,\emptyset}$

computes the possibly true atoms. Therefore, the greatest unfounded set is given by $BASE(P) - \mathrm{lfp}(T_{P_1,\emptyset})$. This set has the property required for $\mathcal{A}$ in Definition 4.11:

Lemma 4.12 (Loop Detection) Let P_1 be a ground program and let

$$P_2 = \{(A \leftarrow \mathcal{B}) \in P_1 \mid A \in \mathrm{lfp}(T_{P_1,\emptyset})\}.$$

Then $P_1 \mapsto_L P_2$ (unless $P_2 = P_1$), and P_2 is irreducible w.r.t. $\mapsto_L$.

PROOF. See Appendix A. □

As shown in [NS96], the set $\mathrm{lfp}(T_{P_1,\emptyset})$ of possible atoms can be computed by the linear time algorithm of [DG84]. As a consequence, also the operation of *loop detection* can be performed in time that is linear in the size of P_1. We compute $\mathrm{lfp}(T_{P_1,\emptyset})$ and delete all rules $A \leftarrow \mathcal{B}$ with $A \notin \mathrm{lfp}(T_{P_1,\emptyset})$.

Example 4.13. Consider again the program *Loop* of Example 4.10. After deleting the second rule by *negative reduction*, *loop detection* can be applied using the unfounded set $\mathcal{A} = \{q, r\}$ to delete the third and fourth rule.

Since more than one application of loop detection may be needed for some programs, in our final rewriting system we have to iterate all five transformations.

Definition 4.14. Let $\mapsto_X$ denote our final rewriting system:

$$\mapsto_X := \mapsto_P \cup \mapsto_N \cup \mapsto_S \cup \mapsto_F \cup \mapsto_L.$$

Thus $\mapsto_X$ can be seen as the extension of Fittings least fixpoint operator Φ_P by *loop detection* (see Theorem 4.9). A normal form w.r.t. the rewriting system $\mapsto_X$ is called a program remainder:

Definition 4.15. (Program Remainder) Let P be a program. Let the program $\widehat{P}$ be a normal form of P w.r.t. $\mapsto_X$. Then $\widehat{P}$ is called a *program remainder* of P.

Lemma 4.16 Every program P has a program remainder, i.e. the rewriting system $\mapsto_X$ is terminating.

PROOF. All our transformations strictly reduce the total number of occurring literals. So if $P_0 := \mathit{ground}(P)$ and P_i is any program with $P_{i-1} \mapsto_X P_i$, $i \geq 1$, then we must reach a program P_n in which no further transformation can be applied. This is a program remainder of P. □

Theorem 4.17 (Computation of Well-Founded Semantics) For every program P and remainder $\widehat{P}$, the well-founded model W_P^* of P satisfies exactly those positive and negative literals which are immediately obvious in $\widehat{P}$, i.e.

$$W_P^* = \mathit{known}_{BASE(P)}(\widehat{P}).$$

PROOF. See Appendix A. □

From this it is also simple to derive an explicit characterization of the program remainder: From the Herbrand instantiation of the program we obtain the program remainder by deleting every rule instance with a body literal which is false in the

well-founded model, and removing from the remaining rule instances the body literals which are true in the well-founded model. This is a kind of GELFOND/LIFSCHITZ-transformation [GL88].

Theorem 4.18 (Unique Program Remainder) Let $\widehat{P}$ be a remainder of a program P. The following holds:

$$\widehat{P} = \{A \leftarrow (\mathcal{B} - W_P^*) \mid A \leftarrow \mathcal{B} \in \mathit{ground}(P) \text{ and for every } B \in \mathcal{B}\colon \sim B \notin W_P^*\}.$$

Thus the remainder of a program is uniquely determined and we will write in the following $\widehat{P}$ for the program remainder of P.

PROOF. See Appendix A. □

Corollary 4.19 The rewriting system $\mapsto_X$ is confluent, strongly terminating and the program remainder $\widehat{P}$ is the unique normal form of $\mathit{ground}(P)$ w.r.t. $\mapsto_X$.

Remember that, as a consequence of unlimited unfolding, the residual program, i.e. the normal form of the rewriting system of [BD98c], can grow to a size that is exponential in the size of the original program. In our approach the program does not grow at all since all transformations strictly reduce the size of the program.

Lemma 4.20 (Quadratic Time complexity) Let P be a ground program. Let n be the size of P, i.e., the number of occurrences of literals in P. Then P can be transformed into its remainder $\widehat{P}$ in time $O(n^2)$.

PROOF. As shown in [NS96] the normal form of P w.r.t. $\mapsto_{PSNF}$ can be computed in time $O(n)$. Also one application of *loop detection* can be performed in time $O(n)$. Since by each application of *loop detection* the size of the program is strictly reduced, *loop detection* can be applied at most n times. Altogether, the remainder $\widehat{P}$ of P is reached in time $O(n^2)$. □

Note, that quadratic time is needed in the worst case if *loop detection* is to be applied n times. Time complexity can be improved up to linear time, if the number of applications of *loop detection* can be reduced, depending on the dependencies in P.

Example 4.21. Consider the following ground program

$$\begin{array}{ll} p(a) & \leftarrow \mathbf{not}\, p(a), \mathbf{not}\, p(b_1). \\ p(b_1) & \leftarrow \mathbf{not}\, p(c_1), \mathbf{not}\, p(b_2). \\ p(b_2) & \leftarrow \mathbf{not}\, p(c_2), \mathbf{not}\, p(b_3). \\ & \vdots \\ p(b_n) & \leftarrow \mathbf{not}\, p(c_n), \mathbf{not}\, p(b_{n+1}). \\ p(c_2). & \end{array}$$

that results from instantiating the program from Example 3.5 using the given base facts and removing obviously true body literals. Note that this grounding operation can be done in linear time w.r.t. n. This program can be transformed by $2n$ applications of the transformations *positive reduction* or *negative reduction* into its

remainder[1]:

$p(b_1).$
$p(b_4).p(b_6).p(b_8).\ldots.p(b_n).$
$p(c_2).$

5. Intelligent Grounding

The transformations defined in Section 4 are applicable only to ground programs. To compute the well-founded model of a non-ground program P we could in a first step generate its Herbrand instantiation $ground(P)$. However, this is a very costly operation since many irrelevant rule instances are produced.

Example 5.1. Consider again the program from Example 3.5. Since it has about $2n$ constants and the first rule has three different variables, its Herbrand instantiation contains $(2n)^3$ ground instances of the first rule. Thus the construction of the Herbrand instantiation would be more expensive than the model computation using the alternating fixpoint procedure (cf. Example 3.5). Instead we would like to construct a ground instantiation containing only relevant rule instances like that of Example 4.21 with only n ground instances of the first rule.

In this section we show how for *range-restricted* programs a set of relevant ground rule instances can be derived efficiently. Together with the magic set transformation which we will discuss in Section 8, this technique can be applied to answer queries against programs that are not range-restricted.

The key idea of the proposed grounding algorithm is to compute the set of all possible facts by ignoring all negative body literals during a bottom-up fixpoint iteration. This corresponds to the computation of the set U_0 of possible atoms in the alternating fixpoint procedure (cf. Definition 3.2). But when a rule instance is applied to derive a new atom, not only the head atom but the complete rule instance is stored. We call the ground instances derived this way *conditional facts*.

Definition 5.2. (Conditional Fact) A *conditional fact* $A \leftarrow \mathcal{C}$ is a ground rule containing possibly both positive and negative delayed literals in the body.

The following immediate consequence operator derives a new conditional fact if for all positive body literals of a ground rule instance conditional facts are given.

Definition 5.3. (Immediate Consequences with Conditional Facts)
Let P be a range-restricted normal program. Given a set S of conditional facts the operator $\bar{T}_P$ computes the following set of conditional facts:

$$\bar{T}_P(S) := \{A \leftarrow \mathcal{B} \in ground(P) \mid pos(\mathcal{B}) \subseteq heads(S)\}.$$

Since $\bar{T}_P$ is monotonic, its least fixpoint $\mathrm{lfp}(\bar{T}_P)$ exists. Obviously, $\mathrm{lfp}(\bar{T}_P)$ is a subset of $ground(P)$. Because all our rules are range-restricted, the condition

[1] We assume that n is even. Otherwise the second line of the remainder changes slightly.

$pos(\mathcal{B}) \subseteq heads(S)$ binds all the variables and we never have to resort to a "blind instantiation". More precisely, it is also clear that $heads\big(\mathrm{lfp}(\bar{T}_P)\big) = \mathrm{lfp}(T_{P,\emptyset})$. The operators are very similar, only $T_{P,\emptyset}$ "forgets" the bodies of the applied rule instances.

The following lemma embeds this kind of "intelligent grounding" into our transformation-based approach and shows that it can be seen as an efficient implementation of a part our transformation system.

Lemma 5.4 Let P be a range-restricted normal program, $\bar{T}_P$ be defined as above, and let $\mapsto_{LF} := \mapsto_L \cup \mapsto_F$, i.e. the combination of *loop detection* and *failure*. Then we have

$$ground(P) \;\mapsto^*_{LF}\; \mathrm{lfp}(\bar{T}_P).$$

Proof. See Appendix B. □

From Lemma 5.4 it follows that $\mathrm{lfp}(\bar{T}_P)$ can be used as a starting point for the transformation approach.

Corollary 5.5 Let P be a range-restricted normal program. Then $\mathrm{lfp}(\bar{T}_P) \mapsto^*_X \widehat{P}$ holds.

Example 5.6. Let P contain the rules and facts from Example 3.5. Then $\mathrm{lfp}(\bar{T}_P)$ contains

$$\begin{array}{ll}
p(a) & \leftarrow t(a,a,b_1), \mathbf{not}\, p(a), \mathbf{not}\, p(b_1).\\
p(b_1) & \leftarrow t(b_1,c_1,b_2), \mathbf{not}\, p(c_1), \mathbf{not}\, p(b_2).\\
p(b_2) & \leftarrow t(b_2,c_2,b_3), \mathbf{not}\, p(c_2), \mathbf{not}\, p(b_3).\\
\multicolumn{2}{c}{\vdots}\\
p(b_n) & \leftarrow t(b_n,c_n,b_{n+1}), \mathbf{not}\, p(c_n), \mathbf{not}\, p(b_{n+1}).\\
p(c_2) & \leftarrow p_0(c_2).
\end{array}$$

together with the base facts.

The grounding algorithm can be made more effective by evaluating a program P step by step according to the SCCs of its static dependency graph. Due to the confluence of our transformation system, we can transform the program P SCC by SCC. The following operator enables us to perform the grounding of P also SCC by SCC. These results are not surprising and we refer to [ZBF97] for more details. We will here present only the main result concerning the intelligent grounding.

Definition 5.7. Let P be a program and R be a ground program. Given a set S of conditional facts, the operator $\hat{T}_{P,R}$ computes the following set of conditional facts:

$$\hat{T}_{P,R}(S) \;:=\; \{A \leftarrow \mathcal{B} \in ground(P) \mid \begin{array}{l} pos(\mathcal{B}) \subseteq heads(R) \cup heads(S) \text{ and} \\ neg(\mathcal{B}) \cap facts(R) = \emptyset\}. \end{array}$$

If R is empty, we obtain the operator $\bar{T}_P$ of Definition 5.3 by $\bar{T}_P(S) = \hat{T}_{P,\emptyset}(S)$.

Since this operator is monotonic, its least fixpoint exists. The following definition describes how it can be used for SCC-oriented program grounding and transforma-

tion.

Definition 5.8. (SCC-Oriented Evaluation) Let P be a program. Let $P^{(1)}, \ldots, P^{(m)}$ be a partition of the set of rules of P such that no predicate occurring in $P^{(i)}$ occurs in the heads of $P^{(j)}$ for any $1 \leq j < i \leq m$. Let the ground programs $R_0, \ldots, R_m$ be defined as follows:

- $R_0 := \emptyset$,
- R_i is the program remainder of $R_{i-1} \cup \mathrm{lfp}(\hat{T}_{P^{(i)}, R_{i-1}})$, for $0 < i \leq m$.

The following theorem states that this evaluation scheme is correct and computes the remainder of P.

Theorem 5.9 (SCC-Oriented Evaluation [ZBF97]) Let P be a program. Let the sequence $R_0, \ldots, R_m$ be defined as above. Then R_m is the remainder of P.

Note, that the bottom-up fixpoint computation of the grounding as described above performs an implicit *loop detection* for each SCC. As a consequence of this, for a large program class no explicit loop detection has to be applied.

Lemma 5.10 (Avoiding Loop Detection [ZBF97]) Let P be a program. Let P contain no predicate which depends on itself both, positively and negatively (i.e. through negation). Then during the SCC-oriented grounding and transformation of P as defined in Definition 5.8 no explicit application of *loop detection* is necessary.

Another optimization is to immediately remove obviously true body literals from the conditional facts already during the fixpoint computation. By this way normal (unconditional) facts are produced whenever the truth value of all body literals is known and conditional facts are produced only if the predicate of a negative body literal is in the same SCC as the head predicate of a rule or a body literal depends already on a conditional fact.

Example 5.11. Consider again the program from Example 3.5. If during the grounding algorithm body literals that are known to be true are immediately removed, the simplified ground program from Example 4.21 is constructed, together with the base facts. Only literals with the recursive predicate p appear in conditions.

Note, that for *predicate stratified programs* the SCC-oriented grounding with immediate removal of true body literals always computes directly the well-founded model of the program. Since the truth values of all negative body literals are defined in lower SCCs, there is no need to produce any conditional facts. Thus the optimized grounding algorithm is a smooth generalization of the standard bottom-up fixpoint computation. Stratified parts of the program are evaluated as usual. Conditional facts are only generated for negative recursion.

6. Relation to Alternating Fixpoint

The computation of the set of all possibly true facts that is needed for the *loop detection* $\mapsto_L$ transformation (cf. Definition 4.11) is very similar to the computa-

tion of possible facts performed in each iteration step of the alternating fixpoint procedure (c.f. Definition 3.2). There is an even closer correspondence as we will show in this section.

Our main result is that a particular ordering of our transformations corresponds exactly to the alternating fixpoint procedure. To characterize the computation performed by the alternating fixpoint procedure we define a sequence $(\bar{K}_i, \bar{U}_i)_{i\geq 0}$ where $\bar{K}_i$ and $\bar{U}_i$ are constructed by the program transformations defined in this paper and that is closely related to the sequence $(K_i, U_i)_{i\geq 0}$ computed by the alternating fixpoint procedure (cf. Definition 3.2).

Definition 6.1. (Alternating Fixpoint by Program Transformations)
Let P be a normal logic program. We define the sequence $(\bar{K}_i, \bar{U}_i)_{i\geq 0}$ of sets $\bar{K}_i$ and $\bar{U}_i$ of conditional facts as follows:

$$\begin{array}{rl} \bar{K}_0 & \text{is the normal form of } \mathit{ground}(P) \text{ w.r.t. } \mapsto_{LFS}, \\ & \text{i.e., } \mathit{ground}(P) \mapsto^*_{LF} \mathrm{lfp}(\bar{T}_P) \mapsto^*_S \bar{K}_0, \\ \bar{U}_0 & \text{is the normal form of } \bar{K}_0 \text{ w.r.t. } \mapsto_{NLF}, \\ i > 0 : \bar{K}_i & \text{is the normal form of } \bar{U}_{i-1} \text{ w.r.t. } \mapsto_{PS}, \\ i > 0 : \bar{U}_i & \text{is the normal form of } \bar{K}_i \text{ w.r.t. } \mapsto_{NLF}. \end{array}$$

The computation terminates when the sequence becomes stationary, i.e., when a fixpoint is reached in the sense that

$$(\bar{K}_j, \bar{U}_j) = (\bar{K}_{j+1}, \bar{U}_{j+1}).$$

Lemma 6.2 (Confluence of $\mapsto_{PS}$ and $\mapsto_{NLF}$) In Definition 6.1, the normal form of $\bar{U}_i$, $i \geq 0$, w.r.t. $\mapsto_{PS}$ and the normal form of $\bar{K}_i$, $i \geq 0$, w.r.t. $\mapsto_{NLF}$ are uniquely determined.

PROOF. See Appendix C. □

Comment 6.6.1. Due to Corollary 4.19 and Lemma 5.4 the sequence $(\bar{K}_i, \bar{U}_i)_{i\geq 0}$ of Definition 6.1 reaches a fixpoint $(\bar{K}_j, \bar{U}_j)$ and computes the program remainder $\widehat{P} = \bar{K}_j = \bar{U}_j$.

The computation of the sets K_i of true facts by the alternating fixpoint procedure corresponds to the application of *success* and *positive reduction* on the transformation side. The computation of the sets U_i of possible facts corresponds to deleting conditional facts by applying *negative reduction*, *loop detection*, and *failure*. The relation between the alternating fixpoint computation and the sequence of program transformations defined above is described more precisely by the following theorem.

Theorem 6.3 Let P be a normal logic program. Let $(K_i, U_i)_{i\geq 0}$ and $(\bar{K}_i, \bar{U}_i)_{i\geq 0}$ be the sequences given by Definitions 3.2 and 6.1, respectively. Then the following holds: $K_i = \mathit{facts}(\bar{K}_i)$ and $U_i = \mathit{heads}(\bar{U}_i)$.

PROOF. See Appendix C. □

Theorem 6.3 shows that the program transformation approach will never do more work than the alternating fixpoint procedure. Although the number of conditional facts in the sets $\bar{K}_i$ and $\bar{U}_i$ may be greater than the number of unconditional facts in the corresponding sets K_i and U_i, the same number of computation steps is needed

to derive them. Several conditional facts having the same head mean that an atom can be derived by more than one rule application. Thus, also the alternating fixpoint computation derives facts more than once, but performs a duplicate elimination afterwards. When a new fact is generated by the transformation approach by more than one transformation application to reduce all body literals of a rule instance, also the alternating fixpoint computation has to consider all body literals before the head of an rule instance can be derived.

However, while having a (sometimes) better time complexity, the transformation approach is expected to have a space complexity worse than the alternating fixpoint approach.

As proposed in [KSS95], the sets K_i do not have to be recomputed in each iteration step, but can be maintained in terms of increments. This optimization is automatically performed by the transformation approach, since conditional facts that were transformed into facts by the deletion of their body literals are never changed afterwards.

Note that the results of this section show that the alternating fixpoint procedure can be seen as one particular implementation of our approach (given by choosing a fixed ordering of applying the transformations). However, by choosing another ordering that delays the application of *loop detection* until no other transformation is applicable, the computation scheme will be much more efficient for many programs. This has already been demonstrated in Examples 3.5 and 4.21. We will further illustrate this point in the next section.

7. Regular Strategy Expressions

In Section 6 we have seen that we can choose the evaluation strategy just by rearranging the order of transformation application. In order to specify a transformation order we introduce the following regular expression syntax.

Definition 7.1. A *regular strategy expression* is defined inductively as follows:

- The single letters P, S, N, F, and L are *atomic* regular strategy expressions.
- If e_1 and e_2 are regular strategy expressions, then $e_1 e_2$ is a regular strategy expression.
- If e is a regular strategy expression, then $(e)^*$ is a regular strategy expression.

Definition 7.2. Let P be a ground program. Let e be a regular strategy expression. P' is the result of the *application of e to P*, $P \overset{e}{\mapsto} P'$, if the following conditions hold:

- If e is an atomic regular strategy expression, then P' is the normal form of P w.r.t. $\mapsto_e$.
- If $e = e_1 e_2$, then there exists a ground program P'' such that $P \overset{e_1}{\mapsto} P''$ and $P'' \overset{e_2}{\mapsto} P'$ hold.
- If $e = (e_1)^*$, then P' results from iteratively applying e_1 to P until no changes occur.

Lemma 7.3 (Unique Result) Let P be a ground program. Let e be a regular strategy expression. Then the result of applying e to P is uniquely determined, i.e., there is only one P' such that $P \overset{e}{\mapsto} P'$ holds.

PROOF. The proof is by induction and based on the fact that the normal form of each base transformation is uniquely determined. □

If we compare strategies, we assume that always the grounding algorithm of Section 5 is prepended such that the execution of the strategies always starts from the set of relevant ground rule instances.

Example 7.4. Following Theorem 4.9, the least fixpoint of the fitting operator corresponds to the normalform $fitt(P)$ of the rewriting system $\mapsto_{PSNF}$. Thus the computation of $fitt(P)$ for a given program P can be expressed by the regular strategy expression $(PSNF)^*$. Since a linear time algorithm for the fitting operator applied to ground programs is given in [NS96], for a ground program P the transformation $P \overset{(PSNF)^*}{\mapsto} fitt(P)$ can be executed in linear time w.r.t. the size of P. Note, that here the order of the four transformations is not relevant for the time complexity.

Example 7.5. Following Definition 6.1 and Theorem 6.3 one iteration of the alternating fixpoint procedure can be described by the expressions PS for the computation of $\bar{K}_i$ and NLF for the computation of $\bar{U}_i$. Thus, the model computation of the alternating fixpoint procedure can be described by the regular strategy expression $(PSNLF)^*$. Since each application of the *loop detection* needs linear time, the transformation $P \overset{(PSNLF)^*}{\mapsto} \widehat{P}$ needs quadratic time w.r.t. the size of the given ground program P. Note, that the expression PS is equivalent to the expression $(PS)^*$, and NLF is equivalent to $(NLF)^*$, i.e., the loop is always executed exactly once. This is shown in the proof of Lemma 6.2.

In the expression $(PSNLF)^*$ of Example 7.5 the expensive *loop detection* is applied in every iteration. In the following optimized strategy *loop detection* is removed from the inner loop and is only applied if no other transformation is applicable.

Example 7.6. A more efficient strategy to compute the remainder of a given program is expressed by the regular strategy expression $((PSNF)^*L)^*$. For programs without positive loops, the transformation $P \overset{((PSNF)^*L)^*}{\mapsto} \widehat{P}$ can be executed in linear time, since the inner loop $(PSNF)^*$ computes the remainder and *loop detection* is applied only once. For programs with positive loops a linear number of iterations of the outer loop may be necessary yielding a quadratic time complexity w.r.t to the size of the given ground program.

Example 7.7. Consider again the program from Example 3.5 and its ground instantiation from Example 4.21. Figure 7.1 illustrates execution times of the AFP

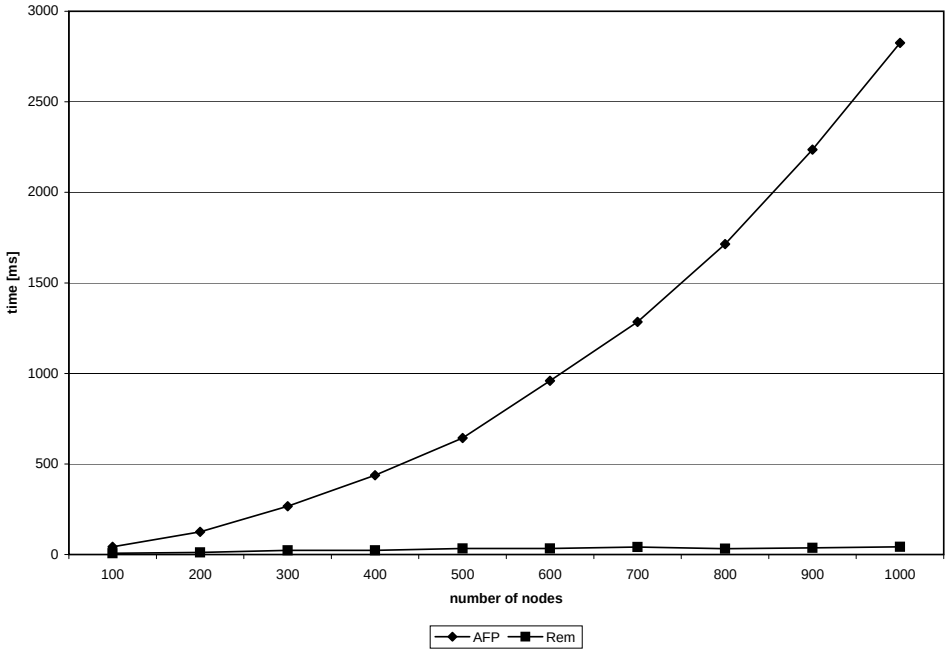

FIGURE 7.1. Experimental Results

strategy from Example 7.5 and the optimized remainder strategy from Example 7.6 for different program sizes n ranging from 100 to 1000. The quadratic and linear character of the two strategies, respectively, can be seen easily.

The tests have been made using our prototypical implementation in java running on a PC with an Intel Pentium III processor with 450 MHz and 128 MB of main memory, MS Windows NT 4.0 operating system and Sun JDK 1.3β java virtual machine. The same test runnning on a Sun Sparc Station with Sun Solaris operating system yields equivalent results.

8. Magic Transformation

It is well-known that the magic set transformation causes problems in the context of the well-founded semantics. The main problem is that not all relevant facts to answer a given query w.r.t. the well-founded model of the original program are derived, if magic facts get undefined truth values in the well-founded model of the magic set transformed program. There are two prominent extensions of the alternating fixpoint procedure that solve this problem by considering undefined magic sets to be true at different steps of the computation. The first approach is the *well-founded magic set* method.

Definition 8.1. (Well-Founded Magic Set [KSS95]) For a given normal program P, any sip strategy S, and a query Q, the *well-founded magic set* method [KSS95]

determines the magic set transformed program $MP = Magic(P, S, Q)$ of P w.r.t. S and Q. Then the set M of true or undefined magic facts in the well-founded model W^*_{MP} is computed using the alternating fixpoint procedure. Finally the well-founded model of $PM = WFMagic(P, S, Q) := MP \cup M$ is computed, again using the alternating fixpoint procedure.

The correctness of this method is stated by the following theorem.

Theorem 8.2 (Well-Founded Magic Set [KSS95]) Let P be a normal logic program. Let S be any sip strategy. Let Q be a query. Let $PM = WFMagic(P, S, Q)$ be as defined above. Then the well-founded models of PM and P agree on Q.

To use this result in our approach we extend our transformation system by the following transformation that allows us to consider undefined magic facts to be true.

Definition 8.3. (Magic Reduction) Let P_1 and P_2 be ground programs. P_2 results from P_1 by *magic reduction* ($P_1 \mapsto_M P_2$) iff there is a rule $A \leftarrow \mathcal{B}$ in P_1 with head predicate p and a positive literal $B \in \mathcal{B}$ with predicate $magic_p$ such that $B \in heads(P_1)$ and $P_2 = \big(P_1 - \{A \leftarrow \mathcal{B}\}\big) \cup \big\{A \leftarrow (\mathcal{B} - \{B\})\big\}$.

Our final transformation system is defined as follows.

Definition 8.4. (MX-Transformation) Let $\mapsto_{MX}$ denote our new rewriting system:

$$\mapsto_{MX} := \mapsto_X \cup \mapsto_M.$$

Note, that this new transformation system is not confluent any more. Depending on the application of the *magic reduction* transformation different strategies for handling undefined magic facts can be realized. However, the following theorem states that any strategy of applying the MX-transformation computes the correct answer for the given query.

Theorem 8.5 (Correctness of the MX-Transformation) Let P be a normal logic program. Let S be any sip strategy. Let Q be a query. let $PM = WFMagic(P, S, Q)$ be as defined above. Let P' be any normal form of PM w.r.t. $\mapsto_{MX}$. Then the well-founded models of P' and P agree on Q.

PROOF. See Appendix D. □

We now extends our regular strategy expression syntax by the new transformation.

Definition 8.6. (Magic Reduction Expression) The regular strategy expression M means the computation of the normalform w.r.t. the transformation *magic reduction*, i.e. the removal of all magic filters from rules with heads with non-magic predicates.

This allows us to embed the well-founded magic set method as a special case of our transformation method.

Definition 8.7. (Well-Founded Magic Sets Strategy) The well-founded magic sets regular strategy expression is defined by $(PSNLF)^*M(PSNLF)^*$.

Lemma 8.8 Let MP be a magic set transformed ground program. The execution of the regular strategy expression $(PSNLF)^*M(PSNLF)^*$ applied to MP corresponds to the computation of the well-founded magic sets method of Definition 8.1.

PROOF. This result follows from Theorem 6.3, Example 7.5, and Definitions 8.3 and 8.6. □

It is obvious that the well-founded magic set strategy can be optimized by delaying the *loop detection* transformation.

Definition 8.9. (Well-Founded Remainder Strategy) The *well-founded remainder* regular strategy expression is defined by $((PSNF)^*L)^*M((PSNF)^*L)^*$.

Lemma 8.10 Let MP be a magic set transformed ground program. The execution of the regular strategy expression $((PSNF)^*L)^*M((PSNF)^*L)^*$ applied to MP is guaranteed to need not more work than the computation corresponding to the expression $(PSNLF)^*M(PSNLF)^*$ and is more efficient for many programs.

PROOF. This statement follows directly from Theorem 6.3. □

One disadvantage of both strategies is that the application of the *magic reduction* transformation all undefined magic facts are considered to be true and removed from all bodies of non-magic rules. This can cause unnecessary computations since as a consequence of making some magic facts true it can happen that other magic facts become false, i.e., the computation of the corresponding subquery is not needed to answer the query. This is the motivation for another approach. The *magic alternating fixpoint* method uses a modified version of the alternating fixpoint procedure where in every computation of true facts all magic facts that were undefined in the preceding computation of true or possible facts are considered true.

Definition 8.11. (Magic Alternating Fixpoint [Mor96]) Let P be a normal logic program. Let S be any sip strategy. Let Q be a query. Let $MP = Magic(P, S, Q)$ be the magic transformed program of P w.r.t. S and Q. Let MP^+ denote the subprogram consisting of the definite rules of MP. The *magic alternating fixpoint procedure* computes the sequence $(K_i, U_i)_{i\geq 0}$ defined as follows:

$$\begin{aligned} K_0 &:= \mathrm{lfp}(T_{MP^+}) \\ U_0 &:= \mathrm{lfp}(T_{MP,K_0}) \\ i > 0 : K_i &:= \mathrm{lfp}(T_{P \cup magic_heads(U_{i-1}),U_{i-1}}) \\ i > 0 : U_i &:= \mathrm{lfp}(T_{P,K_i}). \end{aligned}$$

where $magic_heads(U_{i-1})$ denotes the set of atoms with magic predicates that appear as heads of conditional facts in U_{i-1}.

The correctness of this method is stated by the following theorem.

Theorem 8.12 (Magic Alternating Fixpoint [Mor96]) Let P be a normal logic program. Let S be any sip strategy. Let Q be a query. Let the sequence $(K_i, U_i)_{i\geq 0}$ be defined as above. Then the station gets stationary, i.e., there is an j such that $K_j = K_{j+1}$ and $U_j = U_{j+1}$. Furthermore, the corresponding interpre-

tation $I = K_j \cup \sim(BASE(P) - U_j)$ and the well-founded model W_P^* of P agree on Q.

This approach dynamically changes the set of true magic atoms. The advantage of this method is stated by the following theorem.

Theorem 8.13 (Comparison [KSS95, Mor96]) Let P be a normal logic program. Let S be any sip strategy. Let Q be a query. The magic alternating fixpoint procedure is guaranteed to terminate after no more iterations than the well-founded magic sets approach, and the set of magic facts in the final set U_j of the magic alternating fixpoint procedure is a subset of the set M of true or undefined magic sets in the well-founded magic sets approach.

In the magic alternating fixpoint procedure [Mor96] undefined magic facts are considered to be true only if this enables the derivation of new facts. This can be described by a special case of our *magic reduction* transformation. The *restricted magic reduction* transformation allows to remove a magic filter with undefined truth value from a non-magic rule only if it is the only body literal, i.e., the application transforms the rule into a fact.

Definition 8.14. (Restricted Magic Reduction) Let P_1 and P_2 be ground programs. P_2 results from P_1 by *restricted magic reduction* ($P_1 \mapsto_R P_2$) iff there is a rule $A \leftarrow \mathcal{B}$ in P_1 with head predicate p and *only one* body literal $B \in \mathcal{B}$ with predicate *magic_p* such that $B \in heads(P_1)$ and $P_2 = \big(P_1 - \{A \leftarrow \mathcal{B}\}\big) \cup \big\{A\big\}$.

It is easy to see that this *restricted magic reduction* transformation is sufficient to evaluate magic set transformed program.

Lemma 8.15 (Correctness of Restricted MX-Transformation)
Let MP be a magic set transformed ground program. Let P' a normal form of MP w.r.t. the rewriting system $\mapsto_X \cup \mapsto_R$. Let P'' be a normal form of P' w.r.t. $\mapsto_{MX}$. Then $known_{BASE(P)}(P') = known_{BASE(P)}(P'')$ holds.

PROOF. The only transformation applicable to P' can be $\mapsto_M$ but without producing new facts. Thus no other transformation will become applicable and no new facts can be produced and no rules can be deleted from P'.

The following definition allows us to use the *restricted magic reduction* transformation in regular strategy expressions.

Definition 8.16. (Restricted Magic Reduction Expression) The regular strategy expression R means the computation of the normalform w.r.t. the transformation *restricted magic reduction*, i.e. the removal of all magic filters that are the only body literals of rules with heads with non-magic predicates.

Now we can give express the magic alternating fixpoint procedure as a special case of our rewriting system $\mapsto_{MX}$.

Definition 8.17. (Magic Alternating Fixpoint Strategy) The Magic Alternating Fixpoint regular strategy expression is defined by $(P(SR)^*NLF)^*$.

Lemma 8.18 Let MP be a magic set transformed ground program. The execution

of the regular strategy expression $(P(SR)^*NLF)^*$ applied to MP corresponds to the computation of the magic alternating fixpoint procedure of Definition 8.11.

PROOF. The magic alternating fixpoint method modifies the alternating fixpoint procedure such that during the computation of true atoms undefined magic filters are considered to be true. This means that we have to replace the expression S in the description of the alternating fixpoint procedure by the expression $(SR)^*$ that also computes the normalform of the *success* transformation but additionally is allowed to remove undefined magic filters from non-magic rules if this leads directly to the derivation of facts.

Remark 8.8.1. (Monotonicity) Note, that the sequence of sets K_i of true facts is not monotonically increasing in the magic alternating fixpoint procedure, because some magic facts and facts depending on them that have been considered true once may possibly not be recomputed in a later iteration step. However, in the magic alternating fixpoint transformation, magic conditional facts may be deleted, but non-magic facts once derived as unconditional facts are never deleted afterwards. Thus, our transformational counterpart of the magic alternating fixpoint procedure has the monotonicity property that is missing in the original approach of Morishita.

Examining the magic alternating fixpoint strategy expressions there are two possible optimizations. The more obvious optimization is to delay the application of *loop detection*. The second optimization concerns the time when undefined magic facts are to be considered true. In [KSS95] it is argued that using magic facts too early may cause unnecessary computations since the magic undefined magic facts may become false later. Therefore we delay also the application of the *restricted magic reduction* in the following strategy.

Definition 8.19. (Magic Remainder Strategy) The Magic Remainder regular strategy expression is defined by $(((PSNF)^*R)^*L)^*$.

Lemma 8.20 Let MP be a magic set transformed ground program. The execution of the regular strategy expression $((PSNF)^*R)^*L)^*$ applied to MP is guaranteed to need not more work than the computation corresponding to the expression $(P(SR) * NLF)^*$ and is more efficient for many programs.

PROOF. The expression $((PSNF)^*R)^*$ can be executed in linear time w.r.t. the given ground program. For the expression $(PSNF)^*$ this has been discussed before. For the *restricted magic reduction* each rule has to be checked at most once, i.e., if all body literals except the magic filter have been removed. This operation adds extra costs that are also linear in the size of MP. The only expensive operation is the *loop detection* transformation which is removed from the inner loop and is applied only if no other transformation is applicable.

The following proposition is the main result of this section.

Proposition 8.21 (Magic Remainder) The computation corresponding to the magic remainder strategy is guaranteed to need not work than the computation performed by the well-founded magic sets method or the magic alternating fixpoint and is more efficient for many programs.

The statement of this proposition is illustrated by the following example.

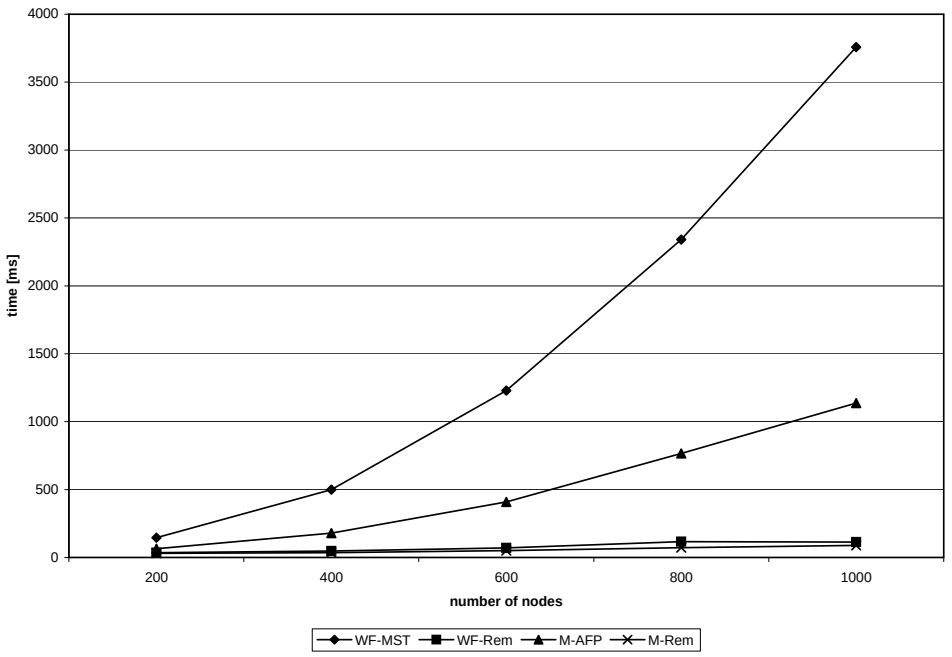

FIGURE 8.1. Comparison without positive Loops

Example 8.22. Consider the following program MP

$$\begin{array}{ll} p(X) & \leftarrow m_p(X), t(X,Y,Z), \mathbf{not}\, p(Y), \mathbf{not}\, p(Z). \\ p(X) & \leftarrow m_p(X), p_0(X). \\ m_p(a). & \\ m_p(Y) & \leftarrow m_p(X), t(X,Y,Z). \\ m_p(Z) & \leftarrow m_p(X), t(X,Y,Z), p(Y). \end{array}$$

that results from the program P from Example 3.5 by the magic set transformation for the query $p(a)$ together with the base facts

$$p_0(c_{\frac{n}{4}}), t(a,a,b), t(b_1,c_1,b_2), t(b_2,c_2,b_3), \ldots, t(b_n,c_n,b_{n+1}).$$

Figure 8.1 illustrates the execution times of the four strategies of Definitions 8.7 (WF-MST), 8.9 (WF-Rem), 8.17 (M-AFP) and 8.19 (M-Rem) for different program sizes n. It can be seen that the WF-MST strategy makes all magic facts true and needs quadratic time to compute the complete well-founded model of the original program. The M-AFP strategy uses only the relevant magic facts and computes only the facts from $p(b_1)$ to $p(b_{\frac{n}{4}})$, but also needs quadratic time. The strategies WF-Rem and M-Rem compute the result in linear time, since the program does not contain positive loops and thus the *loop detection* is not needed.

Consider now the same program but with an additional rule $p(X) \leftarrow p(X)$. Now every p-fact depends positively on itself and a linear number of applications of *loop detection* is needed to compute the result. Figure 8.2 illustrates the execution times

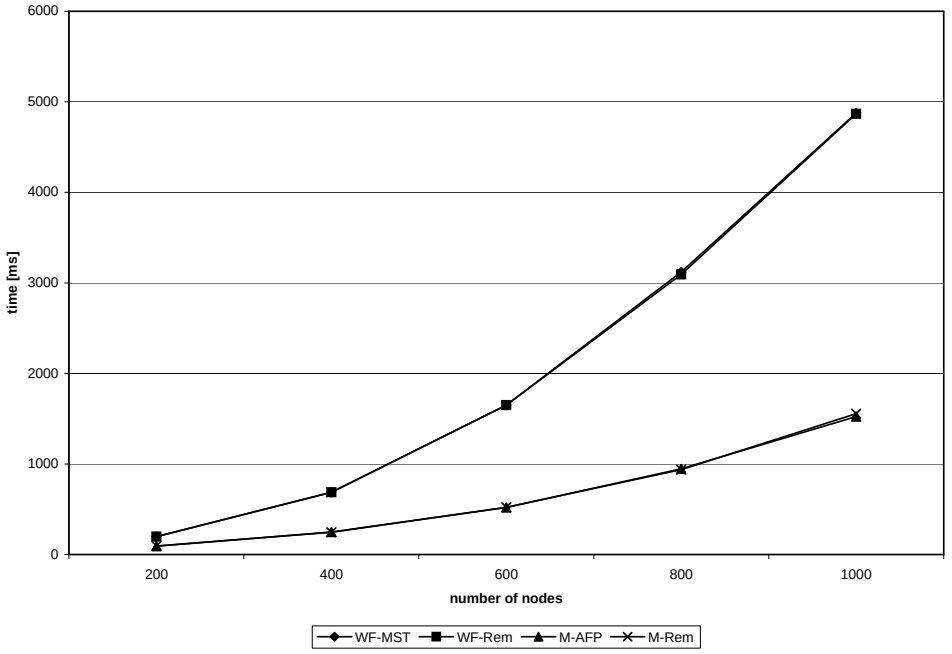

FIGURE 8.2. Comparison with positive Loops

of the four strategies to the modified program. It can be seen the execution times of the strategies WF-MST and M-AFP coincide with the time for the strategies WF-Rem and M-Rem, respectively, and the delay of the *loop detection* has no effect for this program.

9. Discussion

The alternating fixpoint procedure [VG93] (cf. Definition 3.2) is known to have severe inefficiencies. While the sets K_i of true atoms can be computed incrementally the shrinking series of set U_i of possible atoms have to be computed in every iteration step.

The key idea to solve this problem are *conditional facts*. Instead of computing the first set U_0 of possible *atoms*, our intelligent grounding algorithm (cf. Section 5) keeps body literals with unknown truth values in *conditional facts*. By this way a relevant ground instantiation of the original program is derived which can be further processed efficiently by the transformation system defined in Section 4.

The greatest strength of this transformation system is its flexibility. We showed that we can analyze existing evaluation methods and optimize it just by rearranging the order of transformation applications. For example, as stated in Theorem 6.3 and illustrated in Example 7.7 we reduced the quadratic time complexity to a linear just by moving the position of the *loop detection* transformation in the corresponding regular strategy expression (cf. Section 7).

We found valuable implementation hints in [DG84] (a linear time evaluation algorithm for propositional (ground) definite programs, used for the *loop detection* to derive possible atoms ignoring negative body literals), and in [NS96] (a linear time algorithm for the fitting operator, used to computed the normal form of $\mapsto_{PSNF}$). Actually, our optimized remainder strategy $((PSNF)^*L)^*$ is a high level specification of the implementation used by the smodels approach [NS96] to compute the well-founded model of a program.

But our method is not only applicable to range-restricted programs like the alternating fixpoint procedure or the smodels approach, but is also useful for query answering against non range-restricted programs. We showed that our extension of the basic transformation system to magic set transformed programs (cf. Section 8) is superior to the most prominent methods, the well-founded magic sets method [KSS95], and the magic alternating fixpoint procedure [Mor96].

Another prominent bottom-up approach is the *well-founded ordered search* algorithm [SS97]. Like the *ordered search* method [RSS92], it performs a bottom-up fixpoint iteration of a magic set transformed program. But whenever it detects a cyclic negative dependency between ground atoms the alternating fixpoint procedure is applied to evaluate the ground program fragment involved in the negative cycle. Thus, the well-founded ordered search method suffers from the quadratic complexity of the alternating fixpoint procedure and can be optimized by using our transformation approach to resolve negative cycles. The well-founded ordered search method is no special case of our transformation framework since it relies on a more complicated bottom-up computation that keeps track of all dependencies between ground atoms. However, it is similar to our approach in that for stratified programs where no negative cycles occur, it performs a standard bottom-up fixpoint iteration and the overhead for the well-founded semantics is applied only if needed.

....

However, our transformation system can be seen as a subset of the prominent top-down methods like the SLG resolution [CW96] or the tabulated resolution of Bol and Degerstedt [BD98a]. The relationship to SLG is very close, since it can also be characterized as a transformation system, that has to be applied to a more complex data structure containing a system of resolution trees.

While top-down methods are specialized for goal-directed query answering, our method is simple and flexible and can be smoothly integrated with relational bottom-up methods known from database systems...

10. Conclusion

We have presented an algorithm for computing the well-founded semantics of function-free programs which is based on the set of elementary program transformations studied by Brass and Dix [BD94, BD97, BD98b, BD98c]. The time complexity of our algorithm is polynomial in the size of the EDB. If we would unfold all positive body literals in the program remainder (and delete tautologies), we would get the residual program plus some non-minimal conditional facts. So the program remainder can be seen as an efficient encoding of the residual program. It is equivalent to the original program under a wide range of semantics. In fact, our rewriting system $\mapsto_X$ is even weaker than the transformations used

in [BD94, BD97, BD98b, BD98c] because full *unfolding* is not needed to compute the program remainder.

We have shown that our method never does more work than the alternating fixpoint approach. However, there are examples where our method is superior to the alternating fixpoint approach.

Both methods make use of the alternating fixpoint approach. Thus both methods suffer from the inefficient recomputations of possible facts that is inherent to the alternating fixpoint procedure. The authors of the well-founded magic sets method state in [KSS95] that "more research needs to be done to avoid repeated computations".

We have also shown that many bottom-up computation methods that are based on the alternating fixpoint procedure can profit directly from this result. Using conditional facts in the well-founded magic set method [KSS95], the magic alternating fixpoint procedure [Mor96], and the well-founded ordered search algorithm [SS97], unnecessary computations can be avoided in all three approaches. The combination of the goal-directed search strategy of ordered search together with our concepts of conditional facts to store and later reduce possible facts yields a method that is a very competitive bottom-up counterpart of the top-down method SLG [CW96].

REFERENCES

[AMTW99.] K. R. Apt, V. Marek, M. Truszczynski, and D. S. Warren, editors. *The Logic Programming Paradigm: Current Trends and Future Directions*, LNAI, Berlin, 1999. Springer.

[BD94.] Stefan Brass and Jürgen Dix. A disjunctive semantics based on unfolding and bottom-up evaluation. In Bernd Wolfinger, editor, *Innovationen bei Rechen- und Kommunikationssystemen,* (IFIP '94-Congress, Workshop FG2: Disjunctive Logic Programming and Disjunctive Databases), pages 83–91, Berlin, 1994. Springer.

[BD97.] Stefan Brass and Jürgen Dix. Characterizations of the Disjunctive Stable Semantics by Partial Evaluation. *Journal of Logic Programming*, 32(3):207–228, 1997. (Extended abstract appeared in: Characterizations of the Stable Semantics by Partial Evaluation *LPNMR, Proceedings of the Third International Conference, Kentucky*, pages 85–98, 1995. LNCS 928, Springer.).

[BD98a.] Roland N. Bol and Lars Degerstedt. Tabulated resolution for the well-founded semantics. *JLP*, 34(2):67–109, 1998.

[BD98b.] Stefan Brass and Jürgen Dix. Characterizations of the Disjunctive Well-founded Semantics: Confluent Calculi and Iterated GCWA. *Journal of Automated Reasoning*, 20(1):143–165, 1998. (Extended abstract appeared in: Characterizing D-WFS: Confluence and Iterated GCWA. *Logics in Artificial Intelligence, JELIA '96*, pages 268–283, 1996. Springer, LNCS 1126.).

[BD98c.] Stefan Brass and Jürgen Dix. Semantics of (Disjunctive) Logic Programs Based on Partial Evaluation. *Journal of Logic Programming*, accepted for publication, 1998. (Extended abstract appeared in: Disjunctive Semantics Based upon Partial and Bottom-Up Evaluation, *Proceedings of the 12-th International Logic Programming Conference, Tokyo*, pages 199–213, 1995. MIT Press.).

[Bry89.] François Bry. Logic programming as constructivism: A formalization and its application to databases. In *Proc. of the 8th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'89)*, pages 34–50, 1989.

[Bry90.] François Bry. Negation in logic programming: A formalization in constructive logic. In Dimitris Karagiannis, editor, *Information Systems and Artificial Intelligence: Integration Aspects*, number 474 in LNCS, pages 30–46. Springer, 1990.

[BSF95.] Kenneth A. Berman, John S. Schlipf, and John V. Franco. Computing the Well-Founded Semantics Faster. In A. Nerode, W. Marek, and M. Truszczyński, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the Third International Conference*, LNCS 928, pages 113–126, Berlin, June 1995. Springer.

[CSW95.] Weidong Chen, Terrance Swift, and David S. Warren. Efficient Top-Down Computation of Queries under the Well-Founded Semantics. *Journal of Logic Programming*, 24(3):219–245, 1995.

[CW93.] Weidong Chen and David S. Warren. Query-evaluation under the well founded semantics. In *Proc. of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'93)*, pages 168–179, 1993.

[CW96.] Weidong Chen and David S. Warren. Tabled evaluation with delaying for general logic programs. *Journal of the ACM*, 43(1):20–74, 1996.

[DG84.] William F. Dowling and Jean H. Gallier. Linear-time algorithms for testing the satisfiability of propositional horn formulae. *Journal of Logic Programming*, 1(3):267–284, 1984.

[Dix95.] Jürgen Dix. A Classification-Theory of Semantics of Normal Logic Programs: I. Strong Properties. *Fundamenta Informaticae*, XXII(3):227–255, 1995.

[DK89a.] Phan Minh Dung and Kanchana Kanchansut. A fixpoint approach to declarative semantics of logic programs. In *Proc. North American Conference on Logic Programming (NACLP'89)*, pages 604–625, 1989.

[DK89b.] Phan Minh Dung and Kanchana Kanchansut. A natural semantics of logic programs with negation. In *Proc. of the Ninth Conf. on Foundations of Software Technology and Theoretical Computer Science*, pages 70–80, 1989.

[Fit85.] Melvin C. Fitting. A Kripke-Kleene Semantics of logic Programs. *Journal of Logic Programming*, 4:295–312, 1985.

[FSS91.] Burkhard Freitag, Heribert Schütz, and Günther Specht. *LOLA* - a logic language for deductive databases and its implementation. In *Proc. 2nd International Symposium on Database Systems for Advanced Applications (DASFAA '91), Tokyo, Japan, April 2-4, 1991*, pages 216–225, 1991.

[GL88.] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In Robert A. Kowalski and Kenneth A. Bowen, editors, *Logic Programming, Proc. of the 5th Int. Conf. and Symp.*, pages 1070–1080, Cambridge, Mass., 1988. MIT Press.

[KSS95.] David B. Kemp, Divesh Srivastava, and Peter J. Stuckey. Bottom-up evaluation and query optimization of well-founded models. *Theoretical Computer Science*, 146:145–184, 1995.

[Mor96.] Shinichi Morishita. An extension of Van Gelder's alternating fixpoint to magic programs. *Journal of Computer and System Sciences*, 52:506–521, 1996.

[MT91.] Wiktor Marek and Miroslaw Truszczynski. Autoepistemic logic. *Journal of the ACM*, 38(3):588–619, 1991.

[NS96.] Ilka Niemelä and Patrik Simons. Efficient implementation of the well-founded and stable model semantics. Technical Report 7/96, Universität Koblenz-Landau, 1996.

[Prz89.] Teodor Przymusinski. On the declarative and procedural Semantics of logic Programs. *Journal of Automated Reasoning*, 5:167–205, 1989.

[Prz90.] Teodor Przymusinski. Well-founded Semantics Coincides With Three-Valued Stable Semantics. *Fundamenta Informaticae*, XIII:445–463, 1990.

[Prz91.] Teodor Przymusinski. Three-valued non-monotonic formalisms and semantics of logic programs. *Artificial Intelligence*, 49:309–343, 1991. (Extended abstract appeared in: Three-valued non-monotonic formalisms and logic programming. *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89), Toronto, Canada*, pages 341–348, Morgan Kaufmann, 1989.).

[RSS92.] R. Ramakrishnan, D. Srivastava, and S. Sudarshan. Controlling the search in bottom-up evaluation. In Krzysztof Apt, editor, *Proc. Joint International Conference and Symposium on Logic Programming*, pages 273–287, Cambridge, MA, 1992. MIT Press.

[SNV95.] V.S. Subrahmanian, Nau, and Vago. WFS + Branch and Bound = Stable Models. *IEEE Transactions on Knowledge and Data Engineering*, 17(3):362–377, 1995.

[SS97.] Peter J. Stuckey and S. Sudarshan. Well-founded ordered search: Goal-directed bottom-up evaluation of well-founded models. *The Journal of Logic Programming*, 32(3):171–205, 1997.

[SSW94.] Konstantinos Sagonas, Terrance Swift, and David S. Warren. XSB as an efficient deductive database engine. In Richard T. Snodgrass and Marianne Winslett, editors, *Proc. of the 1994 ACM SIGMOD Int. Conf. on Management of Data (SIGMOD'94)*, pages 442–453, 1994.

[VG89.] Allen Van Gelder. The alternating fixpoint of logic programs with negation. In *Proc. of the Eighth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'89)*, pages 1–10, 1989.

[VG93.] Allen Van Gelder. The alternating fixpoint of logic programs with negation. *Journal of Computer and System Sciences*, 47(1):185–221, 1993.

[VGRS91.] Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. The well-founded semantics for general logic programs. *Journal of the Association for Computing Machinary (JACM)*, 38:620–650, 1991.

[ZBF97.] Ulrich Zukowski, Stefan Brass, and Burkhard Freitag. Improving the alternating fixpoint: The transformation approach. In Jürgen Dix, Ulrich Furbach, and Anil Nerode, editors, *Logic Programming and Nonmonotonic Reasoning*, LNAI 1265, pages 40–59, Berlin, 1997. Springer.

[ZF96a.] Ulrich Zukowski and Burkhard Freitag. Adding flexibility to query evaluation for modularly stratified databases. In *Proc. of the 1996 Joint International Conference and Symposium on Logic Programming (JICSLP'96). September 2 – 6, 1996, Bonn, Germany*, pages 304–319. MIT Press, 1996.

[ZF96b.] Ulrich Zukowski and Burkhard Freitag. The differential fixpoint of general logic programs. In Dimitri Boulanger, Ulrich Geske, Fosca Gianotti, and Dietmar Seipel, editors, *Proc. of the Workshop DDLP'96 on Deductive Databases and Logic Programming. 4th Workshop in conjunction with JICSLP'96. Bonn, Germany, September 2 – 6, 1996*, volume 295 of *GMD-Studien*, pages 45–56, St. Augustin, Germany, 1996. GMD.

[ZF97.] Ulrich Zukowski and Burkhard Freitag. The Deductive Database System LOLA. In J. Dix, U. Furbach, and A. Nerode, editors, *Nonmonotonic Extensions of Logic Programming*, LNAI 1265, pages 375–386. Springer, Berlin, 1997.

A. Proofs of Section 4

Definition A.1. (Allowed Transformation) A semantics $\mathcal{S}$ allows a transformation $\mapsto$ iff $\mathcal{S}(P_1) = \mathcal{S}(P_2)$ for all P_1 and P_2 with $P_1 \mapsto P_2$.

Definition A.2. (Deletion of Tautologies) Let P_1 and P_2 be ground programs. Program P_2 results from program P_1 by *deletion of tautologies* ($P_1 \mapsto_T P_2$) iff there is $A \leftarrow \mathcal{B} \in P_1$ such that $A \in \mathcal{B}$ and $P_2 = P_1 - \{A \leftarrow \mathcal{B}\}$.

Definition A.3. (Unfolding) Let P_1 and P_2 be ground programs. Program P_2 results from program P_1 by *unfolding* ($P_1 \mapsto_U P_2$) iff there is a rule $A \leftarrow \mathcal{B}$ in P_1 and a positive literal $B \in \mathcal{B}$ such that

$$P_2 = \big(P_1 - \{A \leftarrow \mathcal{B}\}\big) \cup \big\{A \leftarrow \big((\mathcal{B} - \{B\}) \cup \mathcal{B}'\big) \;\big|\; B \leftarrow \mathcal{B}' \in P_1\big\}.$$

Definition A.4. (Deletion of Non-minimal Rules) Let P_1 and P_2 be ground programs. Program P_2 results from the program P_1 by *deletion of non-minimal rules* ($P_1 \mapsto_M P_2$) iff there are rules $A \leftarrow \mathcal{B}$ and $A \leftarrow \mathcal{B}'$ in P_1 such that $\mathcal{B} \subset \mathcal{B}'$ and $P_2 = P_1 - \{A \leftarrow \mathcal{B}'\}$.

Lemma A.5 If a semantics $\mathcal{S}$ allows *unfolding* and the *deletion of non-minimal rules*, then it also allows *success*.

PROOF. *Success* corresponds to *unfolding*, when the program contains a fact $B \leftarrow \emptyset$ and B is replaced by $\emptyset$. However, there might be further rules about B besides this fact. But then the resulting rules are certainly non-minimal, so we can remove them with $\mapsto_M$. □

Lemma A.6 If a semantics $\mathcal{S}$ allows *unfolding*, then it also allows *failure*.

PROOF. Obviously, the transformation $\mapsto_F$ is a very special case of *unfolding*, i.e., if there are no rules about the unfolded literal. □

Lemma A.7 Let $\mathcal{S}$ be a semantics which allows *unfolding*, *deletion of non-minimal rules*, and *positive* and *negative reduction*. Then also

$$P_1 \mapsto_{PSNF} P_2 \implies \mathcal{S}(P_1) = \mathcal{S}(P_2)$$

for all ground programs P_1 and P_2.

PROOF. The lemma follows immediately from the Lemmata A.5 and A.6. □

PROOF. *(of Lemma 4.9)* Strong termination is obvious, because all transformations strictly reduce the number of literals occurring in the transformed program. To show confluence, we first note that *success* commutes with all other three transformations:

$$\text{If } P \mapsto_S P_1, \text{ and } P_1 \mapsto_{tr} P' \text{ then } P \mapsto_{tr} P_2 \mapsto_S P'$$

and

$$\text{If } P \mapsto_{tr} P_1, \text{ and } P_1 \mapsto_S P' \text{ then } P \mapsto_S P_2 \mapsto_{tr} P'$$

for $tr \in \{$*failure*, *negative reduction*, *positive reduction*$\}$. This commutativity prop-

erty does not hold between the other transformations: consider the program P

$p \leftarrow \mathbf{not}\, q.$
$q.$
$r \leftarrow p.$

If we first apply *negative reduction* to remove the first clause, we can then apply *positive reduction* to simplify the third rule of P. However, this simplification of the third rule is not possible if we start with P.

To show confluence, it suffices (using general results about confluent systems) to prove

If $P \mapsto_X P_1$ and $P \mapsto^*_{P\,N\,F} P_2$ then $P_1 \mapsto^*_{P\,N\,F} P'$ and $P_2 \mapsto^*_{P\,N\,F} P'$,

for $X \in \{\textit{failure}, \textit{negative reduction}, \textit{positive reduction}\} = \{F, N, P\}$. This follows immediately by checking that the following diagrams commute:

$$\begin{array}{ccc} P & \xrightarrow{\mapsto^*_{P\,N\,F}} & P_2 \\ {\scriptstyle P}\downarrow & & \downarrow{\scriptstyle P} \\ P_1 & \xrightarrow[\mapsto^*_{P\,N\,F}]{} & P' \end{array} \qquad \begin{array}{ccc} P & \xrightarrow{\mapsto^*_{P\,N\,F}} & P_2 \\ {\scriptstyle N}\downarrow & & \downarrow{\scriptstyle N} \\ P_1 & \xrightarrow[\mapsto^*_{P\,N\,F}]{} & P' \end{array} \qquad \begin{array}{ccc} P & \xrightarrow{\mapsto^*_{P\,N\,F}} & P_2 \\ {\scriptstyle F}\downarrow & & \downarrow{\scriptstyle F} \\ P_1 & \xrightarrow[\mapsto^*_{P\,N\,F}]{} & P' \end{array}$$

It remains to show the relation to Fitting's least fixpoint of Φ_P. This operator can be best explained if we view it as a mapping between 3-valued interpretations. If we denote the set of truth values by $\mathbf{3}_k$, then $\Phi_P : {\mathbf{3}_k}^{B_P} \longmapsto {\mathbf{3}_k}^{B_P}$; $\mathcal{I} \longmapsto \Phi_P(\mathcal{I})$ is defined as follows: (A is a ground atom)

$$\Phi_P(\mathcal{I})(A) = \begin{cases} \mathbf{t}, & \text{if there is a clause } A \leftarrow L_1, \ldots, L_n \text{ in } P \text{ with:} \\ & \quad — \ \forall i \leq n \text{ we have:} \\ & \qquad * \ \mathcal{I}(L_i) = \mathbf{t}. \\ \mathbf{f}, & \text{if for all clauses } A \leftarrow L_1, \ldots, L_n \in P \text{ we have:} \\ & \quad — \ \exists i \leq n \text{ with:} \\ & \qquad * \ \mathcal{I}(L_i) = \mathbf{f}. \\ \mathbf{u}, & \text{otherwise.} \end{cases}$$

Note that we start with the set $\langle \emptyset, \emptyset \rangle$ and iterate Φ_P until a fixpoint is reached. We will identify $\mathrm{lfp}(\Phi_P)$ with a three-valued model resp. with a set of literals.

To determine the exact relationship between the iterations of Φ_P and the applications of our transformations, we define

$\mathrm{step}_1(P)$:= normal form of P after applying $\mapsto^*_{S\,F}$ and $\mapsto^*_{P\,N}$
$\mathrm{step}_{i+1}(P)$:= normal form of $\mathrm{step}_i(P)$ after applying $\mapsto^*_{S\,F}$ and $\mapsto^*_{P\,N}$

Obviously (by inspection of the definition of Φ_P) we have the following result:

$$\Phi_P \uparrow^i = \textit{known}_{BASE(P)}(\mathrm{step}_i(P)).$$

Clearly, Φ_P reaches a fixpoint if and only if the program P is irreducible with respect to $\mapsto^*_{S\,F\,P\,N}$. The computation of $\mathrm{lfp}(\Phi_P)$ therefore corresponds to a specific ordering of our transformations. □

PROOF. *(of Lemma 4.12)*

1. First we have to show that $\mathcal{A} = BASE(P) - \mathrm{lfp}(T_{P_1,\emptyset})$ satisfies the first requirement in Definition 4.11, namely that for every rule $A \leftarrow \mathcal{B}$ in P_1, if $A \in \mathcal{A}$, then $\mathcal{B} \cap \mathcal{A} \neq \emptyset$. Suppose that this were not the case. Then P_1 would contain a rule $A \leftarrow \mathcal{B}$ with $pos(\mathcal{B}) \subset \mathrm{lfp}(T_{P_1,\emptyset})$ and $A \notin \mathrm{lfp}(T_{P_1,\emptyset})$. But this is impossible, since if $pos(\mathcal{B})$ is already derived by $T_{P_1,\emptyset}$, then A will be derived in the next step.

   So *loop detection* can indeed be applied with this set $\mathcal{A}$, and we get

   $$P_2 = \{(A \leftarrow \mathcal{B}) \in P_1 \mid A \in \mathrm{lfp}(T_{P_1,\emptyset})\}.$$

2. The second part is to show that P_2 is irreducible w.r.t. $\mapsto_L$. Suppose that *loop detection* were applicable to P_2 based on some set $\mathcal{A}'$ of ground atoms. Because of the second and third condition in Definition 4.11, an atom from $\mathcal{A}'$ must appear in at least on rule head of P_2. By the construction of P_2, this atom is contained in $\mathrm{lfp}(T_{P_1,\emptyset})$. But this is impossible: We show by induction on i that $\mathcal{A}' \cap (T_{P_1,\emptyset} \uparrow i) = \emptyset$. For $i = 0$ this is trivial. Let now $A \leftarrow \mathcal{B}$ be applied in the i-th iteration. By the induction hypothesis, we know that $pos(\mathcal{B}) \cap \mathcal{A}' = \emptyset$ and thus $\mathcal{B} \cap \mathcal{A}' = \emptyset$. Since $A \in (T_{P_1,\emptyset} \uparrow i)$, the rule is also contained in P_2. But then the first condition in Definition 4.11 implies that $A \notin \mathcal{A}'$. □

The following lemma shows that *loop detection* is a special case of *deletion of tautologies* and *unfolding*.

Lemma A.8 If $P_1 \mapsto_L P_2$, then also $P_1 \mapsto^*_{TU} P_2$. Thus any semantics $\mathcal{S}$, which allows *unfolding* and *deletion of tautologies*, also allows *loop detection*.

PROOF. Let $\mathcal{A}$ be a set of ground atoms such that for all rules $A \leftarrow \mathcal{B}$ in P_1, if $A \in \mathcal{A}$, then $\mathcal{B} \cap \mathcal{A} \neq \emptyset$. Let $P_2 := \{A \leftarrow \mathcal{B} \in P_1 \mid A \notin \mathcal{A}\}$. We have to show that $P_1 - P_2$, i.e. the rules with an atom from $\mathcal{A}$ in the head, can be deleted by unfolding and elimination of tautologies. We will show $P_1 \mapsto^*_{TU} P_2$ by induction on the number n of atoms from $\mathcal{A}$, which actually appear in rule bodies of $P_1 - P_2$.

If $n = 0$, no atom of $\mathcal{A}$ appears in a body of a rule from $P_1 - P_2$. But since each of the deleted rules must have such an atom in the body, this means $P_1 = P_2$, and thus $P_1 \mapsto^*_{TU} P_2$ trivially holds.

Let now $n > 0$ and $A_0 \in \mathcal{A}$ be an atom which appears in at least one body of a rule from $P_1 - P_2$. Then we first eliminate all tautological rules containing A_0 in head and body. These rules are contained in $P_1 - P_2$, so we are allowed to remove them. After that, we unfold all remaining occurrences of A_0 in bodies of rules which also contain an atom from $\mathcal{A}$ in their head. Let the resulting program be called P_1'. Then we have the following:

- After these operations, A_0 does not appear in any rule body with an atom from $\mathcal{A}$ in the head (Since we have eliminated the tautologies first, the rules about A_0 do not contain A_0 in their bodies.).
- Furthermore, no new atoms appear in the bodies of rules with a head from $\mathcal{A}$ (since $A_0 \in \mathcal{A}$), so the number n has decreased.

- The rules which were deleted by the unfolding steps contain an element of $\mathcal{A}$ in the head, so they are contained in $P_1 - P_2$, and we are really allowed to delete them. The rules which were generated by the unfolding step also contain an atom from $\mathcal{A}$ in the head (since the head is not changed by unfolding). Thus, $P_2 = \{A \leftarrow \mathcal{B} \in P_1' \mid A \notin \mathcal{A}\}$.

The induction hypothesis then gives $P_1' \mapsto_{TU}^* P_2$, and thus together we have $P_1 \mapsto_{TU}^* P_1' \mapsto_{TU}^* P_2$. □

Theorem A.9 (Equivalence of Program Remainder) Let $\mathcal{S}$ be any semantics which allows *unfolding*, *deletion of tautologies*, *positive* and *negative reduction*, and the *deletion of non-minimal rules*. Then

$$\mathcal{S}(\widehat{P}) = \mathcal{S}(P)$$

for every program P with remainder $\widehat{P}$.

PROOF. This follows immediately from the Lemmata A.7 and A.8. □

PROOF. *(of Theorem 4.17)* By Theorem A.9 and since the well-founded semantics allows the transformations of the rewriting system $\mapsto_X$, P and $\widehat{P}$ have the same well-founded model, i.e., $W_P^* = W_{\widehat{P}}^*$. Thus it suffices to show that $W_{\widehat{P}}^* = \mathit{known}_{BASE(P)}(\widehat{P})$. We do this by executing the the alternating fixpoint procedure on $\widehat{P}$.

1. $K_0 = \mathit{facts}(\widehat{P})$: It is clear that at least these facts are contained in K_0, i.e. $\mathit{facts}(\widehat{P}) \subseteq K_0$. To show equality assume that $K_0 - \mathit{facts}(\widehat{P}) \neq \emptyset$. Let $A \leftarrow L_1 \wedge \cdots \wedge L_n$ be the first rule with non-empty body applied by the alternating fixpoint procedure. By Definition 3.2 the L_i are all positive and contained in $\mathit{facts}(\widehat{P})$. But then the *success* transformation is applicable to $\widehat{P}$. This contradicts the irreducibility of $\widehat{P}$.

2. $U_0 = \mathit{heads}(\widehat{P})$: The inclusion $U_0 \subseteq \mathit{heads}(\widehat{P})$ is trivial. Now assume that $U_0 \subset \mathit{heads}(\widehat{P})$ holds. Let $\mathcal{A} := \mathit{BASE}(P) - U_0$ and let $A \leftarrow \mathcal{B}$ be an arbitrary rule of $\widehat{P}$ with $A \in \mathcal{A}$ (such a rule exists because $U_0 \neq \mathit{heads}(\widehat{P})$). If $\mathcal{B} \cap \mathcal{A} \neq \emptyset$, *loop detection* were applicable, contradicting the irreducibility of $\widehat{P}$. Thus we have $\mathcal{B} \cap \mathcal{A} = \emptyset$, and therefore $\mathit{pos}(\mathcal{B}) \subseteq U_0$. But then, since $A \in \mathcal{A}$, i.e., $A \notin U_0$, there must be a negative literal $\mathbf{not}\, B \in \mathcal{B}$ with $B \in K_0$. In this case *negative reduction* would be applicable to $\widehat{P}$ which is impossible because $\widehat{P}$ is irreducible.

3. $K_1 = \mathit{facts}(\widehat{P})$: The inclusion $\mathit{facts}(\widehat{P}) \subseteq K_1$ is trivial, because the fixpoint $K_1 := \mathrm{lfp}(T_{\widehat{P},U_0})$ (cf. Definition 3.2) of course contains all facts of $\widehat{P}$. To show equality assume that $K_1 - \mathit{facts}(\widehat{P}) \neq \emptyset$. Let $A \leftarrow L_1 \wedge \cdots \wedge L_n \in \widehat{P}$ be the first rule with non-empty body applied during the computation of K_1. If it contains a positive body literal L_i, $L_i \in \mathit{facts}(\widehat{P})$ must hold. But then, *success* is applicable, contradicting the irreducibility of $\widehat{P}$. If it contains a negative body literal $L_1 = \mathbf{not}\, B$, $B \notin U_0 = \mathit{heads}(\widehat{P})$ must hold. But then, *positive reduction* is applicable, contradicting again the irreducibility of $\widehat{P}$.

4. Since $K_1 = K_0$, we also have $U_1 = U_0$, so the fixpoint is reached. □

PROOF. *(of Theorem 4.18)* Let us temporarily call the transformation on the right hand side *wred*, i.e. for every program P' let

$$wred(P') := \{A \leftarrow (\mathcal{B} - W^*_{P'}) \mid A \leftarrow \mathcal{B} \in ground(P') \text{ and for every } B \in \mathcal{B}\colon \sim B \notin W^*_{P'}\}.$$

1. We first show that

$$P_1 \mapsto_X P_2 \implies wred(P_1) = wred(P_2).$$

   (a) $P_1 \mapsto_P P_2$: *Positive reduction* removes a negative body literal $\mathbf{not}\, B$ from some rule $A \leftarrow \mathcal{B}$, if $B \notin heads(P_1)$. *Positive reduction* does not change the well-founded model. Furthermore $\mathbf{not}\, B$ is certainly true in the well-founded model. Thus $\mathcal{B} - W^*_{P_1} = (\mathcal{B} - \{\mathbf{not}\, B\}) - W^*_{P_2}$.

   (b) $P_1 \mapsto_N P_2$: *Negative reduction* deletes a rule $A \leftarrow \mathcal{B}$ where $\mathcal{B}$ contains a negative literal $\mathbf{not}\, B$ such that $B \in facts(P_1)$. But then B is true in the well-founded model of P_1 and P_2, so the *wred*-transformation will delete this rule anyway.

   (c) $P_1 \mapsto_S P_2$: *Success* removes a positive body literal B from some rule $A \leftarrow \mathcal{B}$, if $B \in facts(P_1)$. *Success* does not change the well-founded model, and furthermore B is certainly true in this well-founded model. Thus we can conclude $\mathcal{B} - W^*_{P_1} = (\mathcal{B} - \{B\}) - W^*_{P_2}$.

   (d) $P_1 \mapsto_F P_2$: *Failure* deletes a rule $A \leftarrow \mathcal{B}$ where $\mathcal{B}$ contains a positive literal B such that $B \notin heads(P_1)$. But then $\mathbf{not}\, B$ is true in the well-founded model of P_1 and P_2, so the *wred*-transformation will delete this rule anyway.

   (e) $P_1 \mapsto_L P_2$: *Loop detection* w.r.t. to an unfounded set $\mathcal{A}$ deletes all rules $A \leftarrow \mathcal{B}$ where $A \in \mathcal{A}$. By the first condition of Definition 4.11, this also means that $\mathcal{B}$ contains a positive body literal B with $B \in \mathcal{A}$. But since all atoms in the unfounded set $\mathcal{A}$ are false in the well-founded model of P_1 and P_2, the *wred*-transformation will delete these rules anyway.

2. Since $ground(P) \mapsto^*_X \widehat{P}$, we have $wred(ground(P)) = wred(\widehat{P})$. Because the well-founded semantics is instantiation-invariant the definition of *wred* implies $wred(P) = wred(\widehat{P})$. We now only have to show $wred(\widehat{P}) = \widehat{P}$. Assume that $wred(\widehat{P}) \neq \widehat{P}$. Then there is a rule $A \leftarrow \mathcal{B} \in \widehat{P}$ such that there is $B \in \mathcal{B}$ and either $A \leftarrow \mathcal{B}$ has been modified by *wred* because $B \in W^*_{\widehat{P}}$ or entirely deleted because $\sim B \in W^*_{\widehat{P}}$. But we know by Theorem 4.17 that

$$W^*_{\widehat{P}} = known_{BASE(P)}(\widehat{P}).$$

   We can distinguish the following four cases:

   (a) $B \in W^*_{\widehat{P}}$.

      i. B is positive. It follows that $B \in facts(\widehat{P})$. But then *success* is applicable, contradicting the irreducibility of $\widehat{P}$.

      ii. $B = \mathbf{not}\, C$ is negative. It follows that $C \notin heads(\widehat{P})$. But then *positive reduction* is applicable, again contradicting the irreducibility of $\widehat{P}$.

(b) $\sim B \in W^*_{\widehat{P}}$.

i. B is positive. It follows that $\mathbf{not}\, B \in W^*_{\widehat{P}}$ and thus $B \notin \mathit{heads}(\widehat{P})$. But then *failure* is applicable, contradicting the irreducibility of $\widehat{P}$.

ii. $B = \mathbf{not}\, C$ is negative. It follows that $C \in W^*_{\widehat{P}}$ and thus $C \in \mathit{facts}(\widehat{P})$. But then *negative reduction* is applicable, again contradicting the irreducibility of $\widehat{P}$. □

B. Proofs of Section 5

PROOF. *(of Lemma 5.4)* Following Lemma 4.12, $\mathcal{A} := \mathit{BASE}(P) - \mathrm{lfp}(T_{P,\emptyset})$ is a valid reference set for *loop detection*. Let P' be the result of applying *loop detection* w.r.t. this set $\mathcal{A}$ to $\mathit{ground}(P)$. It follows that

$$P' = \{A \leftarrow \mathcal{B} \in \mathit{ground}(P) \mid A \in \mathrm{lfp}(T_{P,\emptyset})\}.$$

Thus $\mathit{heads}(P') \subseteq \mathrm{lfp}(T_{P,\emptyset})$, and we can use *failure* to remove also the rules which contain a positive body literal B with $B \notin \mathrm{lfp}(T_{P,\emptyset})$. Then the result is

$$P'' = \{A \leftarrow \mathcal{B} \in \mathit{ground}(P) \mid \{A\} \cup \mathit{pos}(\mathcal{B}) \subseteq \mathrm{lfp}(T_{P,\emptyset})\}.$$

As mentioned above, it is obvious that $\mathrm{lfp}(T_{P,\emptyset}) = \mathit{heads}(\mathrm{lfp}(\bar{T}_P))$. To prove $P'' = \mathrm{lfp}(\bar{T}_P)$, we show that $P'' \subseteq \mathrm{lfp}(\bar{T}_P)$ and $P'' \supseteq \mathrm{lfp}(\bar{T}_P)$:

1. It is clear that $A \leftarrow \mathcal{B} \in \mathit{ground}(P)$ with

   $$\mathit{pos}(\mathcal{B}) \subseteq \mathrm{lfp}(T_{P,\emptyset}) = \mathit{heads}(\mathrm{lfp}(\bar{T}_P)),$$

   thus $A \leftarrow \mathcal{B}$ must also be contained in $\mathrm{lfp}(\bar{T}_P)$.

2. If $A \leftarrow \mathcal{B} \in \mathit{ground}(P)$ is derived by $\bar{T}_P$, this means that

   $$\mathit{pos}(\mathcal{B}) \subseteq \mathit{heads}(\mathrm{lfp}(\bar{T}_P)) = \mathrm{lfp}(T_{P,\emptyset})$$

   and $A \in \mathit{heads}(\mathrm{lfp}(\bar{T}_P)) = \mathrm{lfp}(T_{P,\emptyset})$.

□

C. Proofs of Section 6

PROOF. *(of Lemma 6.2)* To begin with, we recall Lemma 5.4 and observe that $\bar{K}_0$ is well-defined. The proof by induction is based on the following reasoning:

1. As one possible application of $\mapsto_{PS}$ never interferes with another possible application of $\mapsto_{PS}$, the transformations $\mapsto_P$ and $\mapsto_S$ can be applied in any order and always produce the same result. In particular, it is legal to apply all possible $\mapsto_P$ transformations first followed by all possible $\mapsto_S$ transformations, because $\mapsto_S$ does not derive any new negative information, i.e., it does not delete any rule instance.

2. Analogously, to construct the normal form w.r.t. $\mapsto_{NLF}$ it is legal to apply all possible $\mapsto_N$ transformations first before applying $\mapsto_{LF}$, since *loop detection* and *failure* do not derive any new facts, i.e., they only delete rule instances.

   All applicable $\mapsto_N$ transformations can be applied in any order because they do not interfere with each other and always delete the same set of rule instances.

   Since the greatest unfounded set of a program is the union of all unfounded sets, the normal form w.r.t. $\mapsto_L$ can be constructed by one application of $\mapsto_L$ w.r.t. this greatest unfounded set (cf. Lemma 4.12).

   Finally, rule instances with positive body literals known to be false are deleted by applying $\mapsto_F$. If for an atom A all rule instances having A as their head atom can be deleted by $\mapsto_F$ then $\{A\}$ is an unfounded set and the corresponding rules have already been deleted by $\mapsto_L$. Thus all possible applications of $\mapsto_F$ can be performed in any order.

□

Lemma C.1 Let P_1 be any ground program and let P_2 be the normal form of P_1 w.r.t. $\mapsto_S$ (*success*). Then:

$$P_2 = \{A \leftarrow (\mathcal{B} - facts(P_2)) \wedge \mathbf{not}\,\mathcal{C} \mid A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in P_1\}.$$

PROOF. Every rule in P_2 results from a rule of P_1 by removing positive body literals. Only positive body literals can be removed which are known as facts, and the set of facts can only increase during this process. Since P_2 is irreducible w.r.t. $\mapsto_S$, atoms in $facts(P_2)$ cannot occur any longer as positive body literals. □

Lemma C.2 Let P_1 be any ground program and let P_2 be the normal form of P_1 w.r.t. $\mapsto_N$ (*negative reduction*). Then:

$$P_2 = \{A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in P_1 \mid \mathcal{C} \cap facts(P_1) = \emptyset\}.$$

PROOF. *Negative reduction* deletes the rules which have a negative body literal conflicting with a fact. The set of facts is not changed during this process, i.e. $facts(P_1) = facts(P_2)$. □

Lemma C.3 Let P_1 be any ground program and let P_2 be the normal form of P_1 w.r.t. $\mapsto_P$ (*positive reduction*). Then:

$$P_2 = \{A \leftarrow \mathcal{B} \wedge \mathbf{not}(\mathcal{C} \cap heads(P_1)) \mid A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in P_1\}.$$

PROOF. *Positive reduction* removes the negative body literals for which the corresponding positive atom does not appear as a head literal. The set of head literals is not changed during this process, i.e. $heads(P_1) = heads(P_2)$. □

Lemma C.4 Let P_1 be any ground program and let P_2 be the normal form of P_1 w.r.t. $\mapsto_{LF}$ (*loop detection* and *failure*). Then: $P_2 = \mathrm{lfp}(\bar{T}_{P_1})$.

PROOF.

1. $P_1 \mapsto_{LF}^* \mathrm{lfp}(\bar{T}_{P_1})$ has already been proven in Lemma 5.4.

2. We show that $\mathrm{lfp}(\bar{T}_{P_1})$ is irreducible w.r.t. $\mapsto_{LF}$. Let us assume that $\mapsto_L$ is applicable. Let $\mathcal{A}$ be as in the definition of $\mapsto_L$ (Definition 4.11). Consider the first iteration of the fixpoint iteration in which a rule instance $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$ is derived with $\mathcal{B} \cap \mathcal{A} \neq \emptyset$. There must be such a rule instance since *loop detection* would not be applicable otherwise. Let $B \in \mathcal{B} \cap \mathcal{A}$. Then a rule instance $B \leftarrow \mathcal{B}' \wedge \mathbf{not}\,\mathcal{C}$ has been derived in a previous iteration. This means that $\mathcal{B}' \cap \mathcal{A} = \emptyset$. But this contradicts the first condition in Definition 4.11 which implies that $\mapsto_L$ is not applicable to $\mathrm{lfp}(\bar{T}_{P_1})$. Now let us assume that $\mapsto_F$ is applicable. Then there is a rule instance $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$ and an atom $B \in \mathcal{B}$ that is not contained in $\mathit{heads}(\mathrm{lfp}(\bar{T}_{P_1}))$. But this contradicts the iterative construction of $\mathrm{lfp}(\bar{T}_{P_1})$. □

Lemma C.5 Let P_1 be any ground program and let P_2 be the normal form of P_1 w.r.t. $\mapsto_{LF}$ (*loop detection* and *failure*). Then:

$$P_2 = \big\{A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in P_1 \;\big|\; \mathcal{B} \subseteq \mathit{heads}(P_2)\big\}.$$

Proof. This follows immediately from Lemma C.4: Since P_2 is a fixpoint of $\bar{T}_{P_1}$, every rule $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in P_1$ with $\mathcal{B} \subseteq \mathit{heads}(P_2)$ must be contained in P_2. On the other hand, it follows from the iterative construction of $\mathrm{lfp}(\bar{T}_{P_1})$ that any rule $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$ contained in $\mathrm{lfp}(\bar{T}_{P_1})$ must satisfy $\mathcal{B} \subseteq \mathit{heads}(\mathrm{lfp}(\bar{T}_{P_1}))$. □

Proof. *(of Theorem 6.3)* We do not only show $\mathit{facts}(\bar{K}_i) = K_i$ and $\mathit{heads}(\bar{U}_i) = U_i$ but also the following stronger characterizations of $\bar{K}_i$ and $\bar{U}_i$:

$$\begin{aligned}
\bar{K}_0 &= \big\{A \leftarrow (\mathcal{B} - K_0) \wedge \mathbf{not}\,\mathcal{C} \;\big|\; A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in \mathit{ground}(P),\ \mathcal{B} \subseteq \mathrm{lfp}(T_{P,\emptyset})\big\}\\
\bar{U}_0 &= \big\{A \leftarrow (\mathcal{B} - K_0) \wedge \mathbf{not}\,\mathcal{C} \;\big|\; A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in \mathit{ground}(P),\ \mathcal{B} \subseteq U_0,\ \mathcal{C} \cap K_0 = \emptyset\big\}\\
\bar{K}_i &= \big\{A \leftarrow (\mathcal{B} - K_i) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1}) \;\big|\; A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in \mathit{ground}(P),\ \mathcal{B} \subseteq U_{i-1},\ \mathcal{C} \cap K_{i-1} = \emptyset\big\}, i > 0\\
\bar{U}_i &= \big\{A \leftarrow (\mathcal{B} - K_i) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1}) \;\big|\; A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in \mathit{ground}(P),\ \mathcal{B} \subseteq U_i,\ \mathcal{C} \cap K_i = \emptyset\big\}, i > 0.
\end{aligned}$$

The proof is by induction on i.

1. **Proof for $\bar{K}_0$:** Due to Lemma 5.4 we already know that $\mathit{ground}(P) \mapsto_{LF} \mathrm{lfp}(\bar{T}_P)$. It is easy to show by induction on the number of derivation steps that

   $$\mathrm{lfp}(\bar{T}_P) = \big\{A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \;\big|\; A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in \mathit{ground}(P),\ \mathcal{B} \subseteq T_{P,\emptyset}\big\}$$

   (because $\mathit{heads}(\bar{T}_P \uparrow j) = T_{P,\emptyset} \uparrow j$). By Definition 6.1 $\bar{K}_0$ results from $\mathrm{lfp}(\bar{T}_P)$ by applying $\mapsto_S$ until irreducibility w.r.t. $\mapsto_S$. We show that $\mathit{facts}(\bar{K}_0) = K_0$:

   (a) $K_0 \subseteq \mathit{facts}(\bar{K}_0)$: We show $T_{P^+} \uparrow j \subseteq \mathit{facts}(\bar{K}_0)$ by induction on j: The case $j = 0$ is trivial since $T_{P^+} \uparrow 0 = \emptyset$. Now suppose that A has been derived in step j by applying $A \leftarrow \mathcal{B} \in \mathit{ground}(P^+)$. Then $\mathcal{B} \subseteq T_{P^+} \uparrow (j-1)$ and therefore $\mathcal{B} \subseteq \mathit{facts}(\bar{K}_0)$ by the induction hypothesis. Furthermore we know that $T_{P^+} \uparrow (j-1) \subseteq \mathrm{lfp}(T_{P^+}) \subseteq \mathrm{lfp}(T_{P,\emptyset})$

and therefore $A \leftarrow \mathcal{B} \in \mathrm{lfp}(\bar{T}_P)$. But since $\mathcal{B} \subseteq facts(\bar{K}_0)$ and $\bar{K}_0$ is the normal form of $\mathrm{lfp}(\bar{T}_P)$ w.r.t. $\mapsto_S$, it follows that $A \in facts(\bar{K}_0)$.

(b) $facts(\bar{K}_0) \subseteq K_0$: The proof is by induction on the number of applications of $\mapsto_S$. All facts contained in $\mathrm{lfp}(\bar{T}_P)$ are also contained in $ground(P^+)$ and thus in K_0. Suppose that A is derived from the rule $A \leftarrow \mathcal{B} \in \mathrm{lfp}(\bar{T}_P)$ by j applications of $\mapsto_S$. Then the atoms in $\mathcal{B}$ are derived in fewer steps, and thus $\mathcal{B} \subseteq K_0$ by the induction hypothesis. Since K_0 is a fixpoint of T_{P^+} and $A \leftarrow \mathcal{B} \in ground(P^+)$ it follows that $A \in K_0$.

The characterization of $\bar{K}_0$ now follows from $facts(\bar{K}_0) = K_0$ by Lemma C.1 and the introductory remark on $\mathrm{lfp}(\bar{T}_P)$.

2. **Proof for $\bar{U}_0$:** To construct $\bar{U}_0$ from $\bar{K}_0$ we may compute the normal form of $\bar{K}_0$ w.r.t. $\mapsto_N$ first (see proof of Lemma 6.2) which we denote by $\bar{N}_0$ in the following. By Lemma C.2 and the characterization of $\bar{K}_0$ already shown above

$$\bar{N}_0 = \{A \leftarrow (\mathcal{B} - K_0) \wedge \mathbf{not}\,\mathcal{C} \mid A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in ground(P),\ \mathcal{B} \subseteq \mathrm{lfp}(T_{P,\emptyset}),\ \mathcal{C} \cap K_0 = \emptyset\}$$

holds. Let $\bar{U}_0$ be the normal form of $\bar{N}_0$ w.r.t. $\mapsto_{LF}$. We show that $heads(\bar{U}_0) = U_0$:

(a) $U_0 \subseteq heads(\bar{U}_0)$: Because we have $U_0 = \mathrm{lfp}(T_{P,K_0})$, we prove $T_{P,K_0} \uparrow j \subseteq heads(\bar{U}_0)$ by induction on j. The case $j = 0$ is trivial. Let now $A \in T_{P,K_0} \uparrow j$ have been derived by using the rule instance $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$. Then $\mathcal{B} \subseteq T_{P,K_0} \uparrow (j-1)$ and $\mathcal{C} \cap K_0 = \emptyset$. Also $T_{P,K_0} \uparrow (j-1) \subseteq \mathrm{lfp}(T_{P,K_0}) \subseteq \mathrm{lfp}(T_{P,\emptyset})$, thus $A \leftarrow (\mathcal{B} - K_0) \wedge \mathbf{not}\,\mathcal{C}$ is contained in $\bar{N}_0$. By the induction hypothesis, we have $\mathcal{B} \subseteq heads(\bar{U}_0)$. By Lemma C.5 this implies that $A \leftarrow (\mathcal{B} - K_0) \wedge \mathbf{not}\,\mathcal{C}$ is not deleted by *loop detection* or *failure*. Thus $A \in heads(\bar{U}_0)$.

(b) $heads(\bar{U}_0) \subseteq U_0$: By Lemma C.4, $\bar{U}_0 = \mathrm{lfp}(\bar{T}_{\bar{N}_0})$. Therefore, we show $heads(\bar{T}_{\bar{N}_0} \uparrow j) \subseteq U_0$ by induction on j. The case $j = 0$ is trivial again. Now suppose that $A \leftarrow (\mathcal{B} - K_0) \wedge \mathbf{not}\,\mathcal{C}$ has been derived in step j. Then $(\mathcal{B} - K_0) \subseteq heads(\bar{T}_{\bar{N}_0} \uparrow (j-1))$ and therefore $(\mathcal{B} - K_0) \subseteq U_0$ by the induction hypothesis. But $K_0 \subseteq U_0$ also holds by Lemma 3.4, thus $\mathcal{B} \subseteq U_0$. We also have $\mathcal{C} \cap K_0 = \emptyset$, otherwise the rule instance would not be contained in $\bar{N}_0$. But since $U_0 = \mathrm{lfp}(T_{P,K_0})$, it follows that $A \in U_0$.

Our characterization of $\bar{U}_0$ follows from $heads(\bar{U}_0) = U_0$ by Lemma C.5.

3. **Proof for $\bar{K}_i$, $i \geq 1$:** To construct $\bar{K}_i$, we first apply *positive reduction* to $\bar{U}_{i-1}$ until irreducibility w.r.t. $\mapsto_P$ and denote the resulting set by $\bar{N}_i$. By the induction hypothesis we have $heads(\bar{U}_{i-1}) = U_{i-1}$ and

$$\bar{U}_{i-1} = \{A \leftarrow (\mathcal{B} - K_{i-1}) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-2}) \mid A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in ground(P),\ \mathcal{B} \subseteq U_{i-1},\ \mathcal{C} \cap K_{i-1} = \emptyset\}.$$

By Lemma C.3 and since $U_{i-1} \subseteq U_{i-2}$ (Lemma 3.4) it follows that

$$\bar{N}_i = \{A \leftarrow (\mathcal{B} - K_{i-1}) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1}) \mid A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in ground(P),\ \mathcal{B} \subseteq U_{i-1},\ \mathcal{C} \cap K_{i-1} = \emptyset\}$$

holds. Now we apply $\mapsto_S$ until irreducibility w.r.t. $\mapsto_S$, the result is $\bar{K}_i$. We show that $facts(\bar{K}_i) = K_i$:

(a) $facts(\bar{K}_i) \subseteq K_i$: The proof is by induction on the number of applications of $\mapsto_S$. Let $A \in facts(\bar{K}_i)$. Either $A \in facts(\bar{K}_{i-1})$ and the main induction hypothesis gives $A \in K_{i-1} \subseteq K_i$ or A results from a rule

$$A \leftarrow (\mathcal{B} - K_{i-1}) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1})$$

in $\bar{N}_i$ where $\mathcal{C} \cap U_{i-1} = \emptyset$ and $\mathcal{B} - K_{i-1}$ can be derived as facts with fewer applications of $\mapsto_S$. Then the induction hypothesis gives $(\mathcal{B} - K_{i-1}) \subseteq K_i$. Furthermore $K_{i-1} \subseteq K_i$ by Lemma 3.4. Thus $\mathcal{B} \subseteq K_i$. But then the rule instance $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$ is applicable in the computation of $\mathrm{lfp}(T_{P,U_{i-1}})$ and we get $A \in K_i$.

(b) $K_i \subseteq facts(\bar{K}_i)$: Because of the equality $K_i = \mathrm{lfp}(T_{P,U_{i-1}})$, we show $(T_{P,U_{i-1}} \uparrow j) \subseteq facts(\bar{K}_i)$ by induction on j: The case $j = 0$ is trivial. Now suppose that A has been derived in step j by applying $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$. Then $\mathcal{B} \subseteq \big(T_{P,U_{i-1}} \uparrow (j-1)\big)$ and $\mathcal{C} \cap U_{i-1} = \emptyset$. The first part gives us $\mathcal{B} \subseteq K_i$ which together with Lemma 3.4 implies $\mathcal{B} \subseteq K_i \subseteq U_i \subseteq U_{i-1}$. $\mathcal{C} \cap U_{i-1} = \emptyset$ implies, again using Lemma 3.4, that $\mathcal{C} \cap K_{i-1} = \emptyset$. Together we can conclude that

$$A \leftarrow (\mathcal{B} - K_{i-1}) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1})$$

is contained in $\bar{N}_i$. Moreover, $\mathcal{C} \cap U_{i-1} = \emptyset$ shows that there are actually no negative body literals, so the rule applied is $A \leftarrow (\mathcal{B} - K_{i-1})$. $\mathcal{B} \subseteq \big(T_{P,U_{i-1}} \uparrow (j-1)\big)$ together with the induction hypothesis gives us $\mathcal{B} \subseteq facts(\bar{K}_i)$. But since $\bar{K}_i$ is irreducible w.r.t. $\mapsto_S$, and $A \leftarrow (\mathcal{B} - K_{i-1})$ is contained in $\bar{N}_i$, we can finally conclude $A \in \bar{K}_i$.

Our characterization now follows from $facts(\bar{K}_i) = K_i$ by Lemma C.1 and Lemma 3.4 ($K_{i-1} \subseteq K_i$).

4. **Proof for $\bar{U}_i$, $i \geq 1$:** To construct $\bar{U}_i$ from $\bar{K}_i$, we first compute the normal form of $\bar{K}_i$ w.r.t. $\mapsto_N$ which we denote by $\bar{N}_i$. By Lemma C.2

$$\bar{N}_i = \big\{A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in \bar{K}_i \;\big|\; \mathcal{C} \cap facts(\bar{K}_i) = \emptyset\big\}.$$

As we have just seen, the induction hypothesis implies the characterization of $\bar{K}_i$ and $facts(\bar{K}_i) = K_i$. Therefore and since $K_{i-1} \subseteq K_i$ (see Lemma 3.4) we can conclude that

$$\bar{N}_i = \big\{A \leftarrow (\mathcal{B} - K_i) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1}) \;\big|\; A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C} \in ground(P), \; \mathcal{B} \subseteq U_{i-1},\ \mathcal{C} \cap K_i = \emptyset\big\}$$

holds. Let $\bar{U}_i$ be the normal form of $\bar{N}_i$ w.r.t. $\mapsto_{LF}$. We show $heads(\bar{U}_i) = U_i$:

(a) $heads(\bar{U}_i) \subseteq U_i$: By Lemma C.4, $\bar{U}_i = \mathrm{lfp}(\bar{T}_{\bar{N}_i})$. Therefore, we show $heads(\bar{T}_{\bar{N}_i} \uparrow j) \subseteq U_i$ by induction on j. The case $j = 0$ is trivial. Now suppose that

$$A \leftarrow (\mathcal{B} - K_i) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1})$$

has been derived in step j. Then $(\mathcal{B} - K_i) \subseteq heads(\bar{T}_{\bar{N}_i} \uparrow (j-1))$ and therefore $(\mathcal{B} - K_i) \subseteq U_i$ by the induction hypothesis. But $K_i \subseteq U_i$

also holds by Lemma 3.4, and thus $\mathcal{B} \subseteq U_i$. We also have $\mathcal{C} \cap K_i = \emptyset$, otherwise the derived ground rule would not be contained in $\bar{N}_i$. But $U_i = \mathrm{lfp}(T_{P,K_i})$, and we can apply $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$ in the computation of $\mathrm{lfp}(T_{P,K_i})$, so $A \in U_i$.

(b) $U_i \subseteq \mathit{heads}(\bar{U}_i)$: Since $U_i = \mathrm{lfp}(T_{P,K_i})$, we prove

$$(T_{P,K_i} \uparrow j) \subseteq \mathit{heads}(\bar{U}_i)$$

by induction on j. The case $j = 0$ is trivial. Let now $A \in (T_{P,K_i} \uparrow j)$ be derived by applying the rule instance $A \leftarrow \mathcal{B} \wedge \mathbf{not}\,\mathcal{C}$. Then $\mathcal{B} \subseteq \big(T_{P,K_i} \uparrow (j-1)\big)$ and $\mathcal{C} \cap K_i = \emptyset$. Since $T_{P,K_i} \uparrow (j-1) \subseteq \mathrm{lfp}(T_{P,K_i}) = U_i$ and $U_i \subseteq U_{i-1}$ by Lemma 3.4, we get $\mathcal{B} \subseteq U_{i-1}$. Therefore, the ground rule

$$A \leftarrow (\mathcal{B} - K_i) \wedge \mathbf{not}(\mathcal{C} \cap U_{i-1})$$

is contained in $\bar{N}_i$. From $\mathcal{B} \subseteq \big(T_{P,K_i} \uparrow (j-1)\big)$ the induction hypothesis gives us $\mathcal{B} \subseteq \mathit{heads}(\bar{U}_i)$. But then $A \in \mathit{heads}(\bar{U}_i)$ follows from Lemma C.5.

Our characterization of $\bar{U}_i$ follows from $\mathit{heads}(\bar{U}_i) = U_i$ by Lemma C.5 and Lemma 3.4 ($U_i \subseteq U_{i-1}$). □

D. Proofs of Section 8

PROOF. *(of Lemma 8.5)* This result is based on the fact that all magic facts that are undefined in the well-founded model if MP can be considered true by applications of the *magic success* transformation. In this case the correctness follows from Theorem 8.2. But it can be the case that a conditional magic fact gets deleted before all conditions consisting of its head have been removed by applications of *magic success*. In this case the magic fact has become false because at least one of the body literals in each of its conditional facts have become false. But then also the rule instance for that the magic fact had been generated can be deleted, because its body contains all literals that that the magic fact depends on, due the definition of the magic set transformation. Thus, the deleted magic fact is not needed to answer the given query. As a consequence, all magic facts needed to answer the given query are not deleted and can be considered to be true by applications of the *magic success* transformation. □