An Alternating Well–Founded Semantics for Query Answering in Disjunctive Databases

Dietmar Seipel

University of Würzburg
Am Hubland, D – 97074 Würzburg, Germany
seipel@informatik.uni-wuerzburg.de
Tel. +49 931 888 5026, Fax. +49 931 888 4600

Abstract. The *well–founded semantics* has been introduced for normal databases (i.e. databases that may have default negation in their rule bodies, but do not have disjunctions). In this paper we propose an extension of the well–founded semantics to the disjunctive case. For this purpose we investigate the alternating fixpoint approach of Van Gelder, Ross and Schlipf [16], and develop a suitable generalization to the case of disjunctive rule heads.
Given a disjunctive database $\mathcal{P}$, the new *alternating well–founded semantics* derives a set $ADWFS_{\mathcal{P}}$ of partial Herbrand interpretations of $\mathcal{P}$. This set coincides with the set of minimal models if there are no default negations in the database. For general disjunctive databases it is always not empty (if all rule heads are non–empty), i.e. $ADWFS_{\mathcal{P}}$ is *consistent*. The alternating well–founded semantics is very useful for *query answering* in disjunctive databases. During a fixpoint computation the final set $ADWFS_{\mathcal{P}}$ is approximated by a sequence $(\mathcal{I}_n)_{n \in \mathbb{N}_0}$ of sets $\mathcal{I}_n$ of partial Herbrand interpretations. At any step of the fixpoint computation it holds: If the query already holds in $\mathcal{I}_n$, then the query will also hold in $ADWFS_{\mathcal{P}}$, and the computation can be terminated.
For other semantics like the semantics of stable and partial stable models, so far no computations are known that have this property.

Keywords

disjunctive logic programming, query answering, non–monotonic reasoning, well–founded semantics, handling inconsistency, program transformations

1 Introduction

The alternating fixpoint approach of Van Gelder, Ross and Schlipf [16] has been used for constructing the well–founded semantics of normal databases (without disjunctions). This approach constructs a partial Herbrand interpretation, the so–called *well–founded model* M_W of the database. M_W can be described by two sets of atoms: first, the set T_W of all

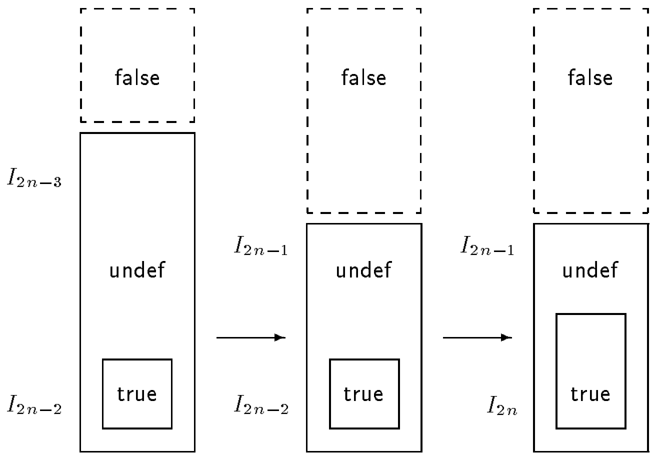

Fig. 1. The Alternating Fixpoint Approach

atoms that are *true* in M_W, and secondly that set F_W of all atoms that are *false* in M_W. Rather than representing the *false* atoms, the construction of M_W deals with the complement set U_W of all atoms that are not *false*, i.e. those atoms that are at least *undefined* in the truth ordering.

T_W and U_W are *approximated* by an alternating sequence $(I_n)_{n \in \mathbb{N}_0}$ of sets of atoms, such that the sets with even (odd) index monotonically approximate T_W (U_W) from below (above), cf. Figure 1. The sets I_n are defined by: $I_0 = \emptyset$, $I_{n+1} = \Gamma_{\mathcal{P}}(I_n)$, for all $n \in \mathbb{N}_0$, where $\Gamma_{\mathcal{P}}(I_n)$ is taken as the minimal Herbrand model of the Gelfond–Lifschitz transformation $\mathcal{P}^{I_n}$ of $\mathcal{P}$ w.r.t. I_n. Then it holds: $I_0 \subseteq \ldots \subseteq I_{2n-2} \subseteq I_{2n} \subseteq \ldots \subseteq T_W \subseteq U_W \subseteq I_{2n-1} \subseteq I_{2n-3} \subseteq \ldots \subseteq I_1$.

Quite frequently, *queries* to a database can be answered without computing the complete sequence $(I_n)_{n \in \mathbb{N}_0}$. For instance for the normal database

$$\mathcal{P} = \{\ c,\ p \leftarrow c \wedge not\, d,\ q \leftarrow c \wedge not\, p,\ a \leftarrow not\, b,\ b \leftarrow not\, a\ \}$$

we get $I_0 = \emptyset$, $I_1 = \{\, a, b, c, p, q \,\}$, and $T_W = I_{2n} = \{\, c, p \,\}$, $U_W = I_{2n+1} = \{\, a, b, c, p \,\}$, for all $n \geq 1$. Already after having computed I_1 it becomes clear that the query "$\mathcal{Q} = \neg d$" evaluates to true (i.e. d is false) under the well–founded semantics, and I_2 shows that the query "$\mathcal{Q} = p$" evaluates to true. In the well–founded model M_W, the atoms "a" and "b" are undefined due to their negative mutual dependency in $\mathcal{P}$.

2 Basic Definitions and Notations

Given a first order language $\mathcal{L}$, a *disjunctive database* $\mathcal{P}$ consists of logical inference rules of the form

$$r = A_1 \vee \ldots \vee A_k \leftarrow B_1 \wedge \ldots \wedge B_m \wedge not\, C_1 \wedge \ldots \wedge not\, C_n, \quad (1)$$

where A_i $(1 \leq i \leq k)$, B_i $(1 \leq i \leq m)$, and C_i $(1 \leq i \leq n)$ are atoms in the language $\mathcal{L}$, $k, m, n \in I\!N_0$, and *not* is the negation–by–default operator.[1] For a rule r it is possible that $m = n = 0$; in this case, r is also called a *fact*. The set of all *ground instances* of the rules in $\mathcal{P}$ is denoted by $gnd\,(\mathcal{P})$. A rule r of the form (1) above is denoted for short as:

$$r = \alpha \leftarrow \beta \wedge not\cdot \gamma, \quad (2)$$

where $\alpha = A_1 \vee \ldots \vee A_k$, $\beta = B_1 \wedge \ldots \wedge B_m$, and $\gamma = C_1 \vee \ldots \vee C_n$.[2] An example for a disjunctive database is given by

$$\mathcal{P} = \{\, a \vee b,\ q \leftarrow b \wedge not\, a,\ p \leftarrow a \wedge not\, b \,\}.$$

Partial Herbrand Interpretations

A *partial (or three–valued) Herbrand interpretation* of a disjunctive database $\mathcal{P}$ is given by a truth function $I : H\!B_{\mathcal{P}} \to \{\, \mathsf{t}, \mathsf{f}, \mathsf{u} \,\}$ that assigns one of the truth values t (true), f (false) or u (undefined) to the atoms of the Herbrand base $H\!B_{\mathcal{P}}$. By $\mathcal{HI}_{\mathcal{P}}$ we denote the set of all partial Herbrand interpretations of $\mathcal{P}$. $I \in \mathcal{HI}_{\mathcal{P}}$ is called a *total Herbrand interpretation*, if all atoms A are mapped to classical truth values t or f. For a partial Herbrand interpretation I we introduce the following notations:

$$I^{\mathsf{t}} = \{\, A^{\mathsf{t}} \mid A \in H\!B_{\mathcal{P}} \wedge I(A) = \mathsf{t} \,\},$$
$$I^{\mathsf{f}} = \{\, A^{\mathsf{f}} \mid A \in H\!B_{\mathcal{P}} \wedge I(A) = \mathsf{f} \,\},$$
$$I^{\mathsf{u}} = \{\, A^{\mathsf{u}} \mid A \in H\!B_{\mathcal{P}} \wedge (\, I(A) = \mathsf{t} \vee I(A) = \mathsf{u} \,) \,\}.$$

Using this notation, there are two ways of representing I as a set of annotated atoms, either by specifying the true and false atoms or by specifying the true and undefined atoms (cf. [14]):

tf–*Representation:* $I^{\mathsf{tf}} = I^{\mathsf{t}} \cup I^{\mathsf{f}}$,
tu–*Representation:* $I^{\mathsf{tu}} = I^{\mathsf{t}} \cup I^{\mathsf{u}}$.

[1] By $I\!N_+ = \{\, 1, 2, 3, \ldots \,\}$ we denote the set of positive natural numbers, whereas $I\!N_0 = \{\, 0, 1, 2, \ldots \,\}$.

[2] Note that γ is a disjunction, and, according to De Morgan's law, $not\cdot \gamma$ is taken to be a conjunction.

Note that in the tu–representation every *true* atom A is recorded as A^{t} and as A^{u}. Note also that the tf–representation is essentially the same as the conventional representation of I as a *set of literals* (where $A^{\mathsf{t}} \mapsto A$ and $A^{\mathsf{f}} \mapsto \neg A$). For a set $\mathcal{I}$ of partial Herbrand interpretations we will use corresponding notations for $\mathsf{v} \in \{ \mathsf{tf}, \mathsf{tu} \}$: $\mathcal{I}^{\mathsf{v}} = \{ I^{\mathsf{v}} \mid I \in \mathcal{I} \}$. By $\mathcal{I} =_{\mathsf{v}} \mathcal{J}$, we denote that $\mathcal{I}^{\mathsf{v}} = \mathcal{J}$. Consider for instance the Herbrand base $H_{B_{\mathcal{P}}} = \{ a, b, c, d, p, q \}$, and the partial Herbrand interpretation I, such that $I(c) = I(p) = \mathsf{t}$, $I(d) = I(q) = \mathsf{f}$, $I(a) = I(b) = \mathsf{u}$. Then I can be represented in the following two ways:

$$I^{\mathsf{tf}} = \{ c^{\mathsf{t}}, p^{\mathsf{t}}, d^{\mathsf{f}}, q^{\mathsf{f}} \}, \quad I^{\mathsf{tu}} = \{ c^{\mathsf{t}}, p^{\mathsf{t}}, c^{\mathsf{u}}, p^{\mathsf{u}}, a^{\mathsf{u}}, b^{\mathsf{u}} \}.$$

Truth Ordering and Knowledge Ordering

There are two common *partial orderings on truth values*, the truth ordering $\leq_{\mathsf{t}}$ and the knowledge ordering $\leq_{\mathsf{k}}$ (cf. [4]):

Truth Ordering:	$\mathsf{f} \leq_{\mathsf{t}} \mathsf{u}$, $\mathsf{u} \leq_{\mathsf{t}} \mathsf{t}$,
Knowledge Ordering:	$\mathsf{u} \leq_{\mathsf{k}} \mathsf{f}$, $\mathsf{u} \leq_{\mathsf{k}} \mathsf{t}$.

These orderings are generalized (pointwise) to partial Herbrand interpretations like follows: for $\mathsf{x} \in \{ \mathsf{t}, \mathsf{k} \}$, $I_1 \leq_{\mathsf{x}} I_2$ iff $I_1(A) \leq_{\mathsf{x}} I_2(A)$, for all $A \in H_{B_{\mathcal{P}}}$. Alternatively, the orderings can be characterized based on the tf–representation and the tu–representation, respectively:

$$I_1 \leq_{\mathsf{k}} I_2 \text{ iff } I_1^{\mathsf{tf}} \subseteq I_2^{\mathsf{tf}}, \text{ and } I_1 \leq_{\mathsf{t}} I_2 \text{ iff } I_1^{\mathsf{tu}} \subseteq I_2^{\mathsf{tu}}.$$

The orderings are generalized to sets of partial Herbrand interpretations like follows: for $\mathsf{x} \in \{ \mathsf{t}, \mathsf{k} \}$,

$$\mathcal{I}_1 \sqsubseteq_{\mathsf{x}} \mathcal{I}_2, \text{ iff for all } I_2 \in \mathcal{I}_2 \text{ there exists } I_1 \in \mathcal{I}_1, \text{ such that } I_1 \leq_{\mathsf{x}} I_2.$$

Finally, for a set $\mathcal{I}$ of partial Herbrand interpretations the minimal elements of $\mathcal{I}$ in the truth ordering and the knowledge ordering, respectively, are denoted by $min_{\mathsf{t}}(\mathcal{I})$ and $min_{\mathsf{k}}(\mathcal{I})$.

Models and Query Answering in Disjunctive Databases

A total Herbrand interpretation I of a disjunctive database $\mathcal{P}$ is a *model* of a ground rule $r = \alpha \leftarrow \beta \wedge not{\cdot}\gamma$, if it holds: if the rule body $\beta \wedge not{\cdot}\gamma$ is true in I, then also the rule head α must be true in I. For a partial Herbrand interpretation I this can be generalized by adding a second

$\wedge$	t	f	u
t	t	f	u
f	f	f	f
u	u	f	u

$\vee$	t	f	u
t	t	t	t
f	t	f	u
u	t	u	u

$\neg$	
t	f
f	t
u	u

Fig. 2. Boolean operations in three–valued logic

condition: if the rule body is undefined in I, then the rule head must be true or undefined. In other words, for a *partial model* of a ground rule r it must hold: $I(\alpha) \geq_{\mathsf{t}} I(\beta) \wedge I(not \cdot \gamma)$, where the Boolean operations on truth values are given in Figure 2.

For model–based semantics, one can compute the set $\mathcal{I}$ of intended models of the database, and one can also derive the set S of formulas that corresponds to $\mathcal{I}$ by means of *sceptical reasoning*. The formulas in S usually are taken from the following sets, which are called the *disjunctive*, *negative*, and the *general Herbrand base*, respectively:

$$\begin{aligned}
\mathit{DHB}_{\mathcal{P}} &= \{ A_1 \vee \ldots \vee A_k \mid A_i \in \mathit{HB}_{\mathcal{P}},\ 1 \leq i \leq k,\ k \in \mathbb{N}_0 \}, \\
\mathit{NHB}_{\mathcal{P}} &= \{ \neg A_1 \vee \ldots \vee \neg A_k \mid A_i \in \mathit{HB}_{\mathcal{P}},\ 1 \leq i \leq k,\ k \in \mathbb{N}_0 \}, \\
\mathit{GHB}_{\mathcal{P}} &= \{ \alpha_d \vee \alpha_n \mid \alpha_d \in \mathit{DHB}_{\mathcal{P}},\ \alpha_n \in \mathit{NHB}_{\mathcal{P}} \}.
\end{aligned}$$

An *Herbrand state* is a set $S \subseteq \mathit{GHB}_{\mathcal{P}}$ of formulas. Sceptical reasoning with a set $\mathcal{I}$ of (partial) Herbrand interpretations consists in deriving the Herbrand state $\mathit{STATE}(S, \mathcal{I})$ of all disjunctions C in a given universe $S \subseteq \mathit{GHB}_{\mathcal{P}}$ which are satisfied by all interpretations in $\mathcal{I}$:

$$\mathit{STATE}(S, \mathcal{I}) = \{ C \in S \mid \forall I \in \mathcal{I} : I \models C \}.$$

A *query* to a disjunctive database $\mathcal{P}$ is given by a disjunction $\mathcal{Q} \in \mathit{GHB}_{\mathcal{P}}$. Given a semantics and the corresponding set $\mathcal{I}$ of intended models of $\mathcal{P}$, then the answer to $\mathcal{Q}$ is yes iff $\mathcal{Q} \in \mathit{STATE}(\mathit{GHB}_{\mathcal{P}}, \mathcal{I})$. For the disjunctive database $\mathcal{P} = \{ a \vee b,\ q \leftarrow b \wedge not\, a,\ p \leftarrow a \wedge not\, b \}$ the set of intended models would usually be the set $\mathcal{I}$ of *perfect models* (cf. [9]) since $\mathcal{P}$ is *stratified* (i.e. there is no recursion through negation in $\mathcal{P}$):

$$\mathcal{I} =_{\mathsf{tf}} \{ \{ a^{\mathsf{t}}, p^{\mathsf{t}}, b^{\mathsf{f}}, q^{\mathsf{f}} \},\ \{ b^{\mathsf{t}}, q^{\mathsf{t}}, a^{\mathsf{f}}, p^{\mathsf{f}} \} \}.$$

Thus, we can derive the disjunctions in the states (subsumed disjunctions are omitted)

$$\begin{aligned}
\mathit{STATE}(\mathit{DHB}_{\mathcal{P}}, \mathcal{I}) &= \{ a \vee b,\ a \vee q,\ b \vee p,\ p \vee q \}, \\
\mathit{STATE}(\mathit{NHB}_{\mathcal{P}}, \mathcal{I}) &= \{ \neg a \vee \neg b,\ \neg a \vee \neg q,\ \neg b \vee \neg p,\ \neg p \vee \neg q \}.
\end{aligned}$$

Besides the "pure" disjunctions in these Herbrand states, the general Herbrand state $\mathit{STATE}(\mathit{GHB}_{\mathcal{P}}, \mathcal{I})$ contains "mixed" disjunctions like $a \vee \neg p$ and $b \vee \neg q$.

3 The Alternating Fixpoint Approach

We will present a new alternating fixpoint approach that uses *annotated constraints* and *annotated databases* for constructing a new type of well–founded models for a disjunctive database $\mathcal{P}$.

The alternating fixpoint approach works on sets $\mathcal{I}$ of partial Herbrand interpretations, and extends them – w.r.t. the knowledge ordering – to other sets $\mathcal{J}$ of partial Herbrand interpretations, such that $\mathcal{I} \sqsubseteq_{\mathsf{k}} \mathcal{J}$. It uses two operators $\Gamma_{\mathcal{P}}^{\downarrow}$ and $\Gamma_{\mathcal{P}}^{\uparrow}$, which take the partial Herbrand interpretations $I \in \mathcal{I}$ and extend them to other partial Herbrand interpretations J such that $I \leq_{\mathsf{k}} J$, cf. Figure 1. The two operators $\Gamma_{\mathcal{P}}^{\downarrow}$ and $\Gamma_{\mathcal{P}}^{\uparrow}$ alternatingly work on the true and the undefined atoms, respectively. For each partial Herbrand interpretation $I \in \mathcal{I}$:

1. the operator $\Gamma_{\mathcal{P}}^{\downarrow}$ computes new partial Herbrand interpretations J, such that $I^{\mathsf{t}} = J^{\mathsf{t}}$ and $I^{\mathsf{u}} \supseteq J^{\mathsf{u}}$, and
2. the operator $\Gamma_{\mathcal{P}}^{\uparrow}$ computes new partial Herbrand interpretations J, such that $I^{\mathsf{u}} = J^{\mathsf{u}}$ and $I^{\mathsf{t}} \subseteq J^{\mathsf{t}}$.

Annotated Constraints

For enforcing the properties given in 1. and 2. above we need the following sets $\mathcal{C}^{\downarrow\mathcal{P},I}$ and $\mathcal{C}^{\uparrow\mathcal{P},I}$ of constraints.

Definition 1. *(Annotated Constraints)*
For a disjunctive database $\mathcal{P}$ and $I \in \mathcal{HI}_{\mathcal{P}}$ we define

$$\mathcal{C}^{\downarrow\mathcal{P},I} = I^{\mathsf{t}} \cup \{\, A^{\mathsf{u}} \leftarrow A^{\mathsf{t}} \mid A \in \mathit{HB}_{\mathcal{P}} \,\} \cup \{\, \leftarrow A^{\mathsf{u}} \mid A \in \mathit{HB}_{\mathcal{P}} \wedge I(A) = \mathsf{f} \,\},$$
$$\mathcal{C}^{\uparrow\mathcal{P},I} = I^{\mathsf{t}} \cup I^{\mathsf{u}} \cup \{\, \leftarrow A^{\mathsf{t}} \mid A \in \mathit{HB}_{\mathcal{P}} \wedge I(A) = \mathsf{f} \,\}.$$

A partial Herbrand interpretation I is extended into another partial Herbrand interpretation J in two phases. When the decreased set J^{u} of undefined atoms is computed, then the constraints $\leftarrow A^{\mathsf{u}}$ forbid the extended interpretation J to contain an annotated atom A^{u} that was not undefined (i.e. that was false) in I. When the increased set J^{t} of true atoms is computed, then the constraints $\leftarrow A^{\mathsf{t}}$ forbid the extended interpretation J to contain any true atom that was not undefined (i.e. that was false) in I.

Annotated Databases

Given a truth value $\mathsf{v} \in \{\, \mathsf{t}, \mathsf{u} \,\}$, for a disjunction $\alpha = A_1 \vee \ldots \vee A_k$ and a conjunction $\beta = B_1 \wedge \ldots \wedge B_m$ of atoms we define $\alpha^{\mathsf{v}} = A_1^{\mathsf{v}} \vee \ldots \vee A_k^{\mathsf{v}}$, $\beta^{\mathsf{v}} = B_1^{\mathsf{v}} \wedge \ldots \wedge B_m^{\mathsf{v}}$. For a disjunctive rule $r = \alpha \leftarrow \beta \wedge not \cdot \gamma$ we define

$$r^{\mathsf{u}} = \alpha^{\mathsf{u}} \leftarrow \beta^{\mathsf{u}} \wedge not \cdot \gamma^{\mathsf{t}},$$
$$r^{\mathsf{t}} = \alpha^{\mathsf{t}} \leftarrow \beta^{\mathsf{t}} \wedge not \cdot \gamma^{\mathsf{u}}.$$

For a disjunctive database $\mathcal{P}$ we define $\mathcal{P}^{\mathsf{u}} = \{\, r^{\mathsf{u}} \mid r \in \mathcal{P} \,\}$, $\mathcal{P}^{\mathsf{t}} = \{\, r^{\mathsf{t}} \mid r \in \mathcal{P} \,\}$, and the *annotated database* $\mathcal{P}^{\mathsf{tu}} = \mathcal{P}^{\mathsf{t}} \cup \mathcal{P}^{\mathsf{u}} \cup \{\, A^{\mathsf{u}} \leftarrow A^{\mathsf{t}} \mid A \in \mathit{HB}_{\mathcal{P}} \,\}$. For instance, the disjunctive database

$$\mathcal{P} = \{\ a \vee b,\ q \leftarrow b \wedge not\, a,\ p \leftarrow a \wedge not\, b\ \},$$

is annotated like follows:

$$\mathcal{P}^{\mathsf{u}} = \{\ a^{\mathsf{u}} \vee b^{\mathsf{u}},\ q^{\mathsf{u}} \leftarrow b^{\mathsf{u}} \wedge not\, a^{\mathsf{t}},\ p^{\mathsf{u}} \leftarrow a^{\mathsf{u}} \wedge not\, b^{\mathsf{t}}\ \},$$
$$\mathcal{P}^{\mathsf{t}} = \{\ a^{\mathsf{t}} \vee b^{\mathsf{t}},\ q^{\mathsf{t}} \leftarrow b^{\mathsf{t}} \wedge not\, a^{\mathsf{u}},\ p^{\mathsf{t}} \leftarrow a^{\mathsf{t}} \wedge not\, b^{\mathsf{u}}\ \}.$$

For any total Herbrand interpretation I of $\mathcal{P}^{\mathsf{tu}}$ we introduce the notation I^{∇} for the partial Herbrand interpretation of $\mathcal{P}$ that is induced by I, i.e. for $A \in \mathit{HB}_{\mathcal{P}}$

$$I^{\nabla}(A) = \begin{cases} \mathsf{t} \text{ if } I(A^{\mathsf{t}}) = \mathsf{t} \\ \mathsf{u} \text{ if } I(A^{\mathsf{u}}) = \mathsf{t} \text{ and } I(A^{\mathsf{t}}) = \mathsf{f} \\ \mathsf{f} \text{ if } I(A^{\mathsf{u}}) = \mathsf{f} \end{cases}$$

For a set $\mathcal{I}$ of total Herbrand interpretations of the annotated database $\mathcal{P}^{\mathsf{tu}}$, let $\mathcal{I}^{\nabla} = \{\, I^{\nabla} \mid I \in \mathcal{I} \,\}$.

Definition 2. *(Gelfond–Lifschitz Transformation)*
Let $\mathcal{P}$ be a disjunctive database $\mathcal{P}$ and $I \in \mathcal{H}\mathit{I}_{\mathcal{P}}$.

(i) For $r = \alpha \leftarrow \beta \wedge not \cdot \gamma \in gnd\,(\mathcal{P})$ we define $r^I = \alpha \leftarrow \beta \wedge \neg I(\gamma)$. The *Gelfond–Lifschitz transformation* of $\mathcal{P}$ is $\mathcal{P}^I = \{\, r^I \mid r \in gnd\,(\mathcal{P}) \,\}$.

(ii) The *Gelfond–Lifschitz transformations* $\mathcal{P}^{\downarrow I}$ and $\mathcal{P}^{\uparrow I}$ are defined by

$$\mathcal{P}^{\downarrow I} = (\mathcal{P}^{\mathsf{u}})^{(I^{\mathsf{t}})},\quad \mathcal{P}^{\uparrow I} = (\mathcal{P}^{\mathsf{t}})^{(I^{\mathsf{u}})}.$$

The Operators $\Gamma_{\mathcal{P}}^{\downarrow}$ and $\Gamma_{\mathcal{P}}^{\uparrow}$

Based on the annotated databases we can define two operators that apply a certain model generation operator to sets of partial Herbrand interpretations, and then *minimize* the resulting sets of partial Herbrand interpretations in the *truth ordering*.

Definition 3. *(The Operators $\Gamma_{\mathcal{P}}^{\downarrow}$ and $\Gamma_{\mathcal{P}}^{\uparrow}$)*
Given a disjunctive database $\mathcal{P}$ and $\mathcal{I} \in 2^{\mathcal{HI}_{\mathcal{P}}}$, the operators

$$\Gamma_{\mathcal{P}}^{\theta} : 2^{\mathcal{HI}_{\mathcal{P}}} \to 2^{\mathcal{HI}_{\mathcal{P}}}$$

for $\theta \in \{ \downarrow, \uparrow \}$ are defined by

$$\begin{aligned} \Gamma_{\mathcal{P}}^{\downarrow}(\mathcal{I}) &= min_{\mathsf{t}}(\bigcup_{I \in \mathcal{I}} \mathcal{MM}(\mathcal{P}^{\downarrow I} \cup \mathcal{C}^{\downarrow \mathcal{P}, I})^{\nabla}), \\ \Gamma_{\mathcal{P}}^{\uparrow}(\mathcal{I}) &= min_{\mathsf{t}}(\bigcup_{I \in \mathcal{I}} \mathcal{MM}(\mathcal{P}^{\uparrow I} \cup \mathcal{C}^{\uparrow \mathcal{P}, I})^{\nabla}). \end{aligned}$$

In the definition above the operator $\mathcal{MM}$ derives a set of total Herbrand interpretations of an annotated database, which is then lifted to a a set of partial Herbrand interpretations of the original database.

Example 4. ($\Gamma_{\mathcal{P}}$–Operators)
For the disjunctive database

$$\mathcal{P} = \{ \, a \vee b, \; q \leftarrow not\, a, \; p \leftarrow not\, b \, \},$$

the *alternating fixpoint approach* derives the following sets of partial Herbrand interpretations, where $\mathcal{I}_0 =_{\mathsf{tu}} \{ A^{\mathsf{u}} \mid A \in H\!B_{\mathcal{P}} \}$, $\mathcal{I}_{2n+1} = \Gamma_{\mathcal{P}}^{\downarrow}(\mathcal{I}_{2n})$, and $\mathcal{I}_{2n+2} = \Gamma_{\mathcal{P}}^{\uparrow}(\mathcal{I}_{2n+1})$, for $n \in I\!N_0$.

$$\begin{aligned} \mathcal{I}_0 &=_{\mathsf{tu}} \{ \{ a^{\mathsf{u}}, b^{\mathsf{u}}, p^{\mathsf{u}}, q^{\mathsf{u}} \} \}, \\ \mathcal{I}_1 &=_{\mathsf{tu}} \{ \{ a^{\mathsf{u}}, p^{\mathsf{u}}, q^{\mathsf{u}} \}, \{ b^{\mathsf{u}}, p^{\mathsf{u}}, q^{\mathsf{u}} \} \}, \\ \mathcal{I}_2 &=_{\mathsf{tu}} \{ \{ a^{\mathsf{t}}, p^{\mathsf{t}}, a^{\mathsf{u}}, p^{\mathsf{u}}, q^{\mathsf{u}} \}, \{ b^{\mathsf{t}}, q^{\mathsf{t}}, b^{\mathsf{u}}, p^{\mathsf{u}}, q^{\mathsf{u}} \} \}, \\ \mathcal{I}_3 &=_{\mathsf{tu}} \{ \{ a^{\mathsf{t}}, p^{\mathsf{t}}, a^{\mathsf{u}}, p^{\mathsf{u}} \}, \{ b^{\mathsf{t}}, q^{\mathsf{t}}, b^{\mathsf{u}}, q^{\mathsf{u}} \} \}, \\ \mathcal{I}_m &= \quad \mathcal{I}_3, \text{ for all } m \geq 3. \end{aligned}$$

It turns out that $\mathcal{I}_3$ is the set of *perfect models* of $\mathcal{P}$.

The disjunctive database $\mathcal{P}$ of the following example is known from literature (cf. [11]). For $\mathcal{P}$ there are no *partial stable models.* With the new alternating fixpoint approach the set of minimal Herbrand models is computed as the semantics of $\mathcal{P}$.

Example 5. *($\Gamma_{\mathcal{P}}$–Operators)*
For the disjunctive database

$$\mathcal{P} = \{\, a \vee b \vee c,\ a \leftarrow not\, b,\ b \leftarrow not\, c,\ c \leftarrow not\, a \,\},$$

the *alternating fixpoint approach* derives the following sets of partial Herbrand interpretations, where $\mathcal{I}_{2n+1} = \Gamma^{\downarrow}_{\mathcal{P}}(\mathcal{I}_{2n})$, and $\mathcal{I}_{2n+2} = \Gamma^{\uparrow}_{\mathcal{P}}(\mathcal{I}_{2n+1})$, for $n \in I\!N_0$.

$$\begin{aligned}
\mathcal{I}_0 &=_{\mathsf{tu}} \{\, \{\, a^{\mathsf{u}}, b^{\mathsf{u}}, c^{\mathsf{u}} \,\} \,\},\\
\mathcal{I}_1 &= \ \mathcal{I}_0,\\
\mathcal{I}_2 &=_{\mathsf{tu}} \{\, \{\, a^{\mathsf{t}}, a^{\mathsf{u}}, b^{\mathsf{u}}, c^{\mathsf{u}} \,\},\ \{\, b^{\mathsf{t}}, a^{\mathsf{u}}, b^{\mathsf{u}}, c^{\mathsf{u}} \,\},\ \{\, c^{\mathsf{t}}, a^{\mathsf{u}}, b^{\mathsf{u}}, c^{\mathsf{u}} \,\} \,\},\\
\mathcal{I}_3 &=_{\mathsf{tu}} \{\, \{\, a^{\mathsf{t}}, a^{\mathsf{u}}, b^{\mathsf{u}} \,\},\ \{\, b^{\mathsf{t}}, b^{\mathsf{u}}, c^{\mathsf{u}} \,\},\ \{\, c^{\mathsf{t}}, a^{\mathsf{u}}, c^{\mathsf{u}} \,\} \,\},\\
\mathcal{I}_4 &=_{\mathsf{tu}} \{\, \{\, a^{\mathsf{t}}, b^{\mathsf{t}}, a^{\mathsf{u}}, b^{\mathsf{u}} \,\},\ \{\, b^{\mathsf{t}}, c^{\mathsf{t}}, b^{\mathsf{u}}, c^{\mathsf{u}} \,\},\ \{\, a^{\mathsf{t}}, c^{\mathsf{t}}, a^{\mathsf{u}}, c^{\mathsf{u}} \,\} \,\},\\
\mathcal{I}_m &= \ \mathcal{I}_4, \text{ for all } m \geq 4.
\end{aligned}$$

It turns out that $\mathcal{I}_4$ is the set of *minimal models* of $\mathcal{P}$.

Note that in general we do not get the minimal models. E.g. in Example 4 we got the set of perfect models of $\mathcal{P}$, and the third minimal model $I =_{\mathsf{tf}} \{\, a^{\mathsf{t}}, b^{\mathsf{t}}, p^{\mathsf{f}}, q^{\mathsf{f}} \,\}$ was not derived.

Naive Alternating Fixpoint Computation

In the following we want to contrast the new alternating fixpoint computation scheme to another computation scheme, that we will call *naive*. The *naive alternating computation scheme* is just based on the Gelfond–Lifschitz transformation, and it does not use the annotated constraints. Furthermore, it works just on single sets of atoms, rather than on partial Herbrand interpretations. The naive computation scheme works as follows: $\mathcal{I}_0 = \{\, \emptyset \,\}$, $\mathcal{I}_{n+1} = \Gamma_{\mathcal{P}}(\mathcal{I}_n)$, for all $n \in I\!N_0$, where

$$\Gamma_{\mathcal{P}}(\mathcal{I}) = min_{\mathsf{t}}\,(\, \bigcup_{I \in \mathcal{I}} \mathcal{MM}(\mathcal{P}^I)\,).$$

Note that for normal databases (i.e. databases without disjunctions) both approaches coincide, and they yield the well–known *well–founded semantics*. When there are no disjunctions in the database, then it is not necessary to explicitly enforce the annotated constraints, since they will be fulfilled automatically.

For disjunctive databases (i.e. databases with disjunctions and negations), however, the following example shows that the naive alternating computation approach may not give us an intuitive interpretation, whereas the alternating fixpoint approach does so.

Example 6. (Naive Alternating Fixpoint Computation)
In the naive alternating computation scheme we get the following results for the previous two example databases.

(i) For the disjunctive database

$$\mathcal{P} = \{\, a \vee b,\ q \leftarrow \mathit{not}\, a,\ p \leftarrow \mathit{not}\, b \,\},$$

we get the following sets of partial Herbrand interpretations, as described in Figure 3:

$$\begin{aligned}
&\mathcal{I}_0 = \{\, \emptyset \,\},\\
&\mathcal{I}_1 = \{\, \{a,p,q\},\ \{b,p,q\} \,\},\\
&\mathcal{I}_2 = \{\, \{a,p\}, \{b,p\}, \{a,q\}, \{b,q\} \,\},\\
&\mathcal{I}_m = \mathcal{I}_2, \text{ for all } m \geq 2.
\end{aligned}$$

Here we do not get the set of perfect models, which we get with the alternating computation scheme. Instead we get two additional non–perfect interpretations $\{b,p\}$ and $\{a,q\}$, which certainly are not intended models of the database.[3]

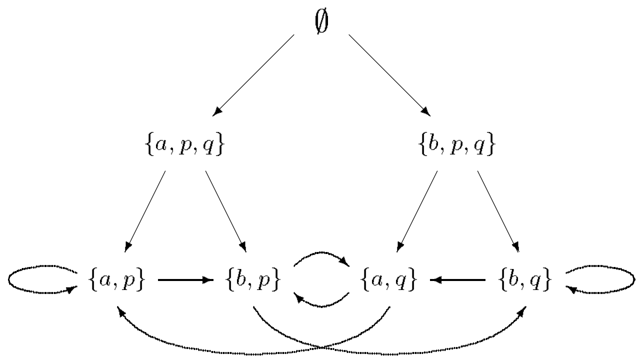

Fig. 3. Naive Computation

(ii) For the disjunctive database

$$\mathcal{P} = \{\, a \vee b \vee c,\ a \leftarrow \mathit{not}\, b,\ b \leftarrow \mathit{not}\, c,\ c \leftarrow \mathit{not}\, a \,\},$$

[3] They are not even models of the database.

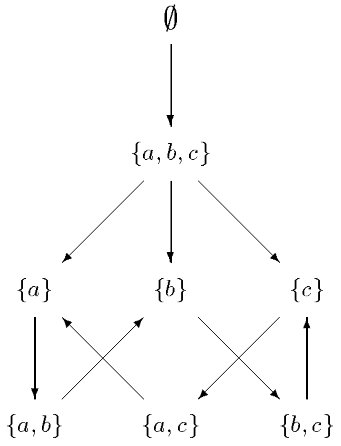

Fig. 4. Naive Computation

we get the following sets of partial Herbrand interpretations, as described in Figure 4:

$$\begin{aligned}
\mathcal{I}_0 &= \{\, \emptyset \,\}, \\
\mathcal{I}_1 &= \{\, \{a, b, c\}, \,\}, \\
\mathcal{I}_2 &= \{\, \{a\}, \{b\}, \{c\} \,\}, \\
\mathcal{I}_3 &= \{\, \{a, b\}, \{b, c\}, \{a, c\} \,\}, \\
\mathcal{I}_{2n} &= \mathcal{I}_2, \text{ for all } n \geq 1, \\
\mathcal{I}_{2n+1} &= \mathcal{I}_3, \text{ for all } n \geq 1.
\end{aligned}$$

Here we do not get the set of minimal models, which we get with the alternating computation scheme. Also the computation does not reach a fixpoint. Instead it loops between the set $\mathcal{I}_3$ of minimal Herbrand models and a set $\mathcal{I}_2$ of three Herbrand interpretations $\{a\}$, $\{b\}$ and $\{c\}$, which are not even models of $\mathcal{P}$. In this approach the meaning of the derived set of partial Herbrand interpretations is not clear at all.

4 The Alternating Well–Founded Semantics *ADWFS*

The composite application of the two operators $\Gamma_{\mathcal{P}}^{\downarrow}$ and $\Gamma_{\mathcal{P}}^{\uparrow}$ defines a new operator $F_{\mathcal{P}}$. The ordinal powers of the operator $F_{\mathcal{P}}$ correspond to a sequence of alternating applications of the operators $\Gamma_{\mathcal{P}}^{\downarrow}$ and $\Gamma_{\mathcal{P}}^{\uparrow}$.

Definition 7. *(The Operator $F_{\mathcal{P}}$)*
Given a disjunctive database $\mathcal{P}$ and $\mathcal{I} \in 2^{\mathcal{HI}_{\mathcal{P}}}$.

(i) The operator

$$F_{\mathcal{P}} : 2^{\mathcal{HI}_{\mathcal{P}}} \to 2^{\mathcal{HI}_{\mathcal{P}}}$$

operates on sets $\mathcal{I} \in 2^{\mathcal{HI}_{\mathcal{P}}}$ of partial Herbrand interpretations:

$$F_{\mathcal{P}}(\mathcal{I}) = \Gamma_{\mathcal{P}}^{\uparrow}(\, \Gamma_{\mathcal{P}}^{\downarrow}(\mathcal{I})\,).$$

(ii) The *ordinal powers* $F_{\mathcal{P}}^{\uparrow\alpha}(\mathcal{I})$ are

$$\begin{aligned}
&F_{\mathcal{P}}^{\uparrow 0}(\mathcal{I}) = \mathcal{I},\\
&F_{\mathcal{P}}^{\uparrow\alpha}(\mathcal{I}) = F_{\mathcal{P}}(F_{\mathcal{P}}^{\uparrow\alpha-1}(\mathcal{I})), \text{ for a successor ordinal } \alpha,\\
&F_{\mathcal{P}}^{\uparrow\alpha}(\mathcal{I}) = lub_k\{\, F_{\mathcal{P}}^{\uparrow\beta}(\mathcal{I}) \mid \beta < \alpha\,\}, \text{ for a limit ordinal } \alpha,
\end{aligned}$$

where $lub_k(\mathcal{X})$ is the *least upper bound* of $\mathcal{X} \subseteq 2^{\mathcal{HI}_{\mathcal{P}}}$ in the knowledge ordering $\sqsubseteq_{\mathsf{k}}$.

(iii) The *ordinal powers* are $F_{\mathcal{P}} \uparrow \alpha = F_{\mathcal{P}}^{\uparrow\alpha}(\bot_F)$, where

$$\bot_F =_{\mathsf{tu}} \{\, A^{\mathsf{u}} \mid A \in H B_{\mathcal{P}}\,\}$$

is the least element of $2^{\mathcal{HI}_{\mathcal{P}}}$ in the knowledge ordering $\sqsubseteq_{\mathsf{k}}$.

Example 8. (Disjunctive Well–Founded Models)
For the disjunctive database

$$\mathcal{P} = \{\, a \vee b,\ q \leftarrow not\, a,\ p \leftarrow not\, b\,\},$$

the *alternating fixpoint approach* derives the following sets of partial Herbrand interpretations:

$$\begin{aligned}
F_{\mathcal{P}} \uparrow 0 &=_{\mathsf{tu}} \{\, \{\, a^{\mathsf{u}}, b^{\mathsf{u}}, p^{\mathsf{u}}, q^{\mathsf{u}}\,\}\,\},\\
F_{\mathcal{P}} \uparrow 1 &=_{\mathsf{tu}} \{\, \{\, a^{\mathsf{t}}, p^{\mathsf{t}}, a^{\mathsf{u}}, p^{\mathsf{u}}, q^{\mathsf{u}}\,\},\ \{\, b^{\mathsf{t}}, q^{\mathsf{t}}, b^{\mathsf{u}}, p^{\mathsf{u}}, p^{\mathsf{u}}\,\}\,\},\\
F_{\mathcal{P}} \uparrow 2 &=_{\mathsf{tu}} \{\, \{\, a^{\mathsf{t}}, p^{\mathsf{t}}, a^{\mathsf{u}}, p^{\mathsf{u}}\,\},\ \{\, b^{\mathsf{t}}, q^{\mathsf{t}}, b^{\mathsf{u}}, p^{\mathsf{u}}\,\}\,\},\\
F_{\mathcal{P}} \uparrow n &= \quad F_{\mathcal{P}} \uparrow 2, \text{ for all } n \geq 2.
\end{aligned}$$

$F_{\mathcal{P}} \uparrow 2$ is the set of *perfect models* of $\mathcal{P}$.

Monotonicity and Fixpoints of the Operators

For the operators $\Gamma_{\mathcal{P}}^{\downarrow}$ and $\Gamma_{\mathcal{P}}^{\uparrow}$ and a set $\mathcal{I}$ of partial Herbrand interpretations it holds $\mathcal{I} \sqsubseteq_{\mathsf{k}} \Gamma_{\mathcal{P}}^{\downarrow}(\mathcal{I})$ and $\mathcal{I} \sqsubseteq_{\mathsf{k}} \Gamma_{\mathcal{P}}^{\uparrow}(\mathcal{I})$. Thus, we also get

$$\mathcal{I} \sqsubseteq_{\mathsf{k}} \Gamma_{\mathcal{P}}^{\uparrow}(\Gamma_{\mathcal{P}}^{\downarrow}(\mathcal{I})) = F_{\mathcal{P}}(\mathcal{I}).$$

By transfinite induction it can be shown that the ordinal powers $F_{\mathcal{P}}^{\uparrow\alpha}(\mathcal{I})$ form a chain in the knowledge ordering:

$$\alpha \leq \beta \Longrightarrow F_{\mathcal{P}}^{\uparrow\alpha}(\mathcal{I}) \sqsubseteq_{\mathsf{k}} F_{\mathcal{P}}^{\uparrow\beta}(\mathcal{I}).$$

From this we can conclude that the operator $F_{\mathcal{P}}$ is *monotonic* on its ordinal powers. Thus, the ordinal powers $F_{\mathcal{P}} \uparrow \alpha$ converge to a fixpoint[4] $F_{\mathcal{P}} \uparrow \gamma$, where γ is called the *closure ordinal* of $F_{\mathcal{P}}$.

Definition 9. *(The Alternating Disjunctive Well–Founded Semantics)*
Given a disjunctive database $\mathcal{P}$, the *alternating disjunctive well–founded semantics* of $\mathcal{P}$ is given by the set

$$\textit{ADWFS}_{\mathcal{P}} = F_{\mathcal{P}} \uparrow \gamma,$$

where γ is the closure ordinal of $F_{\mathcal{P}}$. The partial Herbrand interpretations in $\textit{ADWFS}_{\mathcal{P}}$ are called alternating disjunctive well–founded models.

Consistency

It can be shown that the *alternating fixpoint approach* will always derive a non–empty set of partial models for all disjunctive databases $\mathcal{P}$ without denial rules (i.e. rules with an empty rule head).

Theorem 10. *(Consistency of ADWFS)*
Let $\mathcal{P}$ be a disjunctive database, and let $I_0 =_{\mathsf{tf}} \emptyset$. If $\mathcal{P}^{I_0}$ is logically consistent[5], then $\textit{ADWFS}_{\mathcal{P}} \neq \emptyset$.

The semantics of *stable* and *partial stable models* (cf. [5,6,11]) are among the most prominent semantics for disjunctive databases. Unfortunately, there are databases which are logically consistent, but are inconsistent w.r.t. these semantics. For large databases, small inconsistent parts can prohibit the existence of (partial) stable models.

[4] Note that this fixpoint is the least fixpoint of $F_{\mathcal{P}}$ when restricted to its ordinal powers, but it need not be the least fixpoint of $F_{\mathcal{P}}$ in general, since $F_{\mathcal{P}}$ is not monotonic in general.

[5] I.e. $\mathcal{P}^{\emptyset}$ it has an Herbrand model, which is guaranteed if there are no denial rules in the database $\mathcal{P}$.

5 Conclusions

The disjunctive well–founded semantics has got quite a number of nice properties. If the databases are *normal*, then it is identical with the *well–founded* semantics. For *positive–disjunctive* databases (i.e. databases without default negation) it is identical with the *minimal model* semantics, and for *disjunctive* databases without denial rules there always exist well–founded models, i.e. the semantics is *consistent*.

The computation of alternating disjunctive well–founded models has been implemented within the system DisLog for efficient reasoning in disjunctive databases, cf. [12].

In [14] a characterization of the *partial stable models* (cf. [11]) of a disjunctive database $\mathcal{P}$ in terms of the *total stable models* (cf. [6]) of the annotated database $\mathcal{P}^{\mathsf{tu}}$ is given. Since the construction of $\mathit{ADWFS}_{\mathcal{P}}$ is essentially based on $\mathcal{P}^{\mathsf{tu}}$, too, currently we are trying to investigate the relationship between the two different approaches.

References

1. *S. Brass, J. Dix:* Characterizations of the Disjunctive Stable Semantics by Partial Evaluation, Proc. Third Intl. Conf. on Logic Programming an Non–Monotonic Reasoning (LPNMR'95), Springer LNAI 928, 1995, pp. 85–98, and: Journal of Logic Programming, vol. 32(3), 1997, pp. 207–228.
2. *T. Eiter, N. Leone, D. Sacca:* On the Partial Semantics for Disjunctive Deductive Databases, Annals of Mathematics and Artificial Intelligence, to appear.
3. *B.A. Davey, H.A. Priestley:* Introduction to Lattices and Order, Cambridge University Press, 1990.
4. *M. Fitting:* Bilattices and the Semantics of Logic Programs, Journal of Logic Programming, vol. 11, 1991, pp. 91–116.
5. *M. Gelfond, V. Lifschitz:* The Stable Model Semantics for Logic Programming, Proc. Fifth Intl. Conference and Symposium on Logic Programming (ICSLP'88), MIT Press, 1988, pp. 1070–1080.
6. *M. Gelfond, V. Lifschitz:* Classical Negation in Logic Programs and Disjunctive Databases, New Generation Computing, vol. 9, 1991, pp. 365–385.
7. *J.W. Lloyd:* Foundations of Logic Programming, second edition, Springer, 1987.
8. *J. Lobo, J. Minker, A. Rajasekar:* Foundations of Disjunctive Logic Programming, MIT Press, 1992.
9. *T.C. Przymusinski*: On the declarative semantics of Deductive Databases and Logic Programming, in: Foundations of Deductive Databases and Logic Programming, J. Minker ed., 1988, pp. 193–216.
10. *T.C. Przymusinski*: Extended Stable Semantics for Normal and Disjunctive Programs, Proc. Seventh Intl. Conference on Logic Programming, 1990, pp. 459–477.
11. *T.C. Przymusinski*: Stable Semantics for Disjunctive Programs, New Generation Computing, vol. 9, 1991, pp. 401–424.

12. *D. Seipel:* DisLog – A Disjunctive Deductive Database Prototype, Proc. Twelfth Workshop on Logic Programming (WLP'97), 1997, pp. 136–143.
DisLog is available on the WWW at
"http://www-info1.informatik.uni-wuerzburg.de/databases/DisLog".
13. *D. Seipel, J. Minker, C. Ruiz:* Model Generation and State Generation for Disjunctive Logic Programs, Journal of Logic Programming, vol. 32(1), 1997, pp. 48–69.
14. *D. Seipel, J. Minker, C. Ruiz:* A Characterization of Partial Stable Models for Disjunctive Deductive Databases, Proc. Intl. Logic Programming Symposium (ILPS'97), MIT Press, 1997, pp. 245–259.
15. *D. Seipel:* Partial Evidential Stable Models For Disjunctive Databases, Proc. Workshop on Logic Programming and Knowledge Representation (LPKR'97) at the International Symposium on Logic Programming 1997 (ILPS'97), 1997.
16. *A. Van Gelder, K.A. Ross, J.S. Schlipf:*, Unfounded Sets and Well–Founded Semantics for General Logic Programs, Proc. Seventh ACM Symposium on Principles of Database Systems, 1988, pp. 221–230.
17. *C. Witteveen, G. Brewka:* Skeptical Reason Maintenance and Belief Revision, Journal of Artificial Intelligence, vol. 61, 1993, pp. 1–36.
18. *J.H. You, L.Y. Yuan:* Three-Valued Formalisms of Logic Programming: Is It Needed ?, Proc. Ninth ACM Symposium on Principles of Database Systems (PODS'90), 1990, pp. 172–182.