

Abduction in Well-Founded Semantics and Generalized Stable Models via Tabled Dual Programs

José Júlio Alferes*

Luís Moniz Pereira[†]

Terrance Swift[‡]

June 23, 2003

Abstract

Abductive logic programming offers a formalism to declaratively express and solve problems in areas such as diagnosis, planning, belief revision and hypothetical reasoning. Tabled logic programming offers a computational mechanism that provides a level of declarativity superior to that of Prolog, and which has supported successful applications in fields such as parsing, program analysis, and model checking. In this paper we show how to use tabled logic programming to evaluate queries to abductive frameworks with integrity constraints when these frameworks contain both default and explicit negation. The result is the ability to compute abduction over well-founded semantics with explicit negation and answer sets. Our approach consists of a transformation and an evaluation method. The transformation adjoins to each objective literal O in a program, an objective literal $not(O)$ along with rules that ensure that $not(O)$ will be true if and only if O is false. We call the resulting program a *dual* program. The evaluation method, ABDUAL, then operates on the dual program. ABDUAL is sound and complete for evaluating queries to abductive frameworks whose entailment method is based on either the well-founded semantics with explicit negation, or on answer sets. Further, ABDUAL is asymptotically as efficient as any known method for either class of problems. In addition, when abduction is not desired, ABDUAL operating on a dual program provides a novel tabling method for evaluating queries to ground extended programs whose complexity and termination properties are similar to those of the best tabling methods for the well-founded semantics. A publicly available meta-interpreter has been developed for ABDUAL using the XSB system.

1 Introduction

Abductive logic programming (see e.g. [21]) is a general non-monotonic formalism whose potential for applications is striking. As is well known, problems in domains such as diagnosis, planning, and temporal reasoning can be naturally modeled through abduction. In this paper (which is an extended and revised version with proofs of [2]), we lay the basis for efficiently computing queries over

*A.I. Centre, Faculdade de Ciências e Tecnologia, Univ. Nova de Lisboa, 2825-516 Caparica, Portugal. jjal@di.fct.unl.pt

[†]A.I. Centre, Faculdade de Ciências e Tecnologia, Univ. Nova de Lisboa, 2825-516 Caparica, Portugal. gal.lmp@di.fct.unl.pt

[‡]Department of Computer Science, University of Maryland, College Park, MD, USA. tsswift@cs.umd.edu

ground three-valued abductive frameworks that are based on extended logic programs with integrity constraints, and whose notion of entailment rests on the well-founded semantics and its partial stable models. Both the generalized stable models semantics [22] and the answer set semantics [17] are also captured, their two-valuedness being imposed by means of appropriate integrity constraints.

Our query processing technique, termed ABDUAL, relies on a mixture of program transformation and tabled evaluation. In our abductive framework, a transformation removes default negative literals from both the program over which abduction is to be performed and from the integrity rules. Specifically a *dual transformation* is used, that defines for each objective literal O and its set of rules R , a dual set of rules whose conclusion $\text{not}(O)$ is true if and only if O is false by R . Tabled evaluation of the resulting program turns out to be much simpler than for the original program, whenever abduction over negation is needed. At the same time, termination and complexity properties of tabled evaluation of extended programs are preserved by the transformation when abduction is not needed.

Regarding tabled evaluation, ABDUAL is in the line of SLG evaluation [5] which computes queries to normal programs according to the well-founded semantics. In fact, its definition is inspired by a simplification, for ground programs, of SLG as reformulated in [33]. To it, ABDUAL tabled evaluation adds mechanisms to handle abduction¹, and to deal with the dual programs².

The contributions of this paper are:

- We describe ABDUAL fully and first consider its use over abductive frameworks whose entailment method is based on the well-founded semantics with explicit negation. ABDUAL is sound, complete, and terminating for queries to such frameworks over finite ground programs and integrity rules. Furthermore, ABDUAL is ideally sound and complete for countably infinite ground programs.
- We show that over abductive frameworks whose entailment method is based on the well-founded semantics with explicit negation, the complexity of ABDUAL is in line with the best known methods. In addition, for normal and extended programs — viewed as abductive frameworks containing no abducibles or integrity constraints — query evaluation has polynomial data complexity.
- We provide a transformation that allows ABDUAL to compute generalized stable models and answer sets under a credulous semantics, and show that ABDUAL provides a sound and complete evaluation method for computing such models. Furthermore, the efficiency of ABDUAL in computing generalized stable models is in line with the best known methods.
- Finally, we provide access to an ABDUAL meta-interpreter, written using the XSB system, illustrating how to evaluate ABDUAL in practice and describe how ABDUAL can be applied to medical diagnosis [16], to reasoning about actions [1], and to model-based diagnosis of an electric power grid [20].

¹Namely, by adding abductive contexts to goals, by modifying operations on forests to deal with such contexts, and by having a new operation to abduce literals.

²Namely by introducing a CO-UNFOUNDED SET REMOVAL operation.

2 Preliminaries

2.1 Terminology and assumptions

Throughout this paper, we use the terminology of Logic Programming as defined in, e.g. [25], with the following modifications. An *objective literal* is either an atom A , or the *explicit negation* of A , denoted $\neg A$. If an objective literal O is an atom A , the explicit conjugate of O ($\text{conj}_E(O)$) is the atom $\neg A$; otherwise if O has the form $\neg A$, the explicit conjugate of O is A . A *literal* either has the form O , where O is an objective literal, or $\text{not}(O)$ the *default negation* of O . In the first form, where a literal is simply an objective literal, it is called a *positive literal*; in the second, where it is of the form $\text{not}(O)$, it is called a *negative literal*. Default conjugates are defined similarly to explicit conjugates: the default conjugate ($\text{conj}_D(O)$) of an objective literal O is $\text{not}(O)$, and the default conjugate of $\text{not}(O)$ is O . Thus, every atom is an objective literal and every objective literal is a literal. A *program* P (sometimes also called an extended program), formed over some countable language of function and predicate symbols \mathcal{L}_P , is a countable set of rules of the form $H :- \text{Body}$ in which H is an objective literal, and Body is a possibly empty finite sequence of literals. If no objective literals in a program P contain the explicit negation symbol, P is called *normal*. In either case, the closure of the set of literals occurring in P under explicit and default conjugation is termed $\text{literals}(P)$, while the closure of the set of objective literals occurring in P closed under explicit conjugation is termed $\text{objective_literals}(P)$.

By a *three-valued interpretation* \mathcal{I} of a ground program P we mean a subset of $\text{literals}(P)$. We denote as \mathcal{I}_T the set of objective literals in \mathcal{I} , and as \mathcal{I}_F the set of literals of the form $\text{not}(O)$ in \mathcal{I} . For a ground objective literal, O , if neither O nor $\text{not}(O)$ is in \mathcal{I} , the truth value of O is undefined. An interpretation \mathcal{I} is *consistent* if there is no objective literal O such that $O \in \mathcal{I}_T$ and $\text{not}(O) \in \mathcal{I}_F$; \mathcal{I} is *coherent* if $O \in \mathcal{I}_T$ implies $\text{not}(\text{conj}_E(O)) \in \mathcal{I}_F$ ³. The *information ordering* of interpretations is defined as follows. Given two interpretations, \mathcal{I} and \mathcal{J} , $\mathcal{I} \subseteq_{\text{Info}} \mathcal{J}$ if \mathcal{I}_F is a subset of \mathcal{J}_F , and \mathcal{I}_T is a subset of \mathcal{J}_T . Given an interpretation \mathcal{I} and a set of objective literals \mathcal{S} , $\mathcal{I}|_{\mathcal{S}}$, the restriction of \mathcal{I} to \mathcal{S} , is $\{L | L \in \mathcal{I} \text{ and } (L \in \mathcal{S} \text{ or } (L = \text{not}(O) \text{ and } O \in \mathcal{S}))\}$. Any consistent three-valued interpretation can be viewed as a function from $\text{literals}(P)$ to the set $\{\mathbf{f}, \mathbf{u}, \mathbf{t}\}$. Accordingly, for convenience we assume that the symbols \mathbf{t} and $\text{not}(\mathbf{f})$ belong to every model, while neither \mathbf{u} nor $\text{not}(\mathbf{u})$ belong to any model. For simplicity of presentation, we assume a left-to-right literal selection strategy throughout this paper, although any of the results presented here will hold for any fixed literal selection strategy. Finally, because dual programs (introduced below) allow any literal as the head of a rule, the terms *goal*, *query* and *literal* are used interchangeably.

2.2 The Well-Founded Semantics for Extended Programs

We first recall definitions of the well-founded and stable models for extended programs. The well-founded model can be seen as a double iterated fixed point whose inner operators determine a set of true and false literals at each step.

³A coherent interpretation ensures that if some objective literal is explicitly false (resp. true) then it also must be false (resp. true) by default.

Definition 2.1 For a ground program P , interpretation \mathcal{I} of P and sets \mathcal{O}_1 and \mathcal{O}_2 of ground objective literals

- $Tx_{\mathcal{I}}^P(\mathcal{O}_1) = \{O : \text{there is a clause } O :- L_1, \dots, L_n \in P \text{ and for each } i, 1 \leq i \leq n, L_i \in \mathcal{I} \text{ or } L_i \in \mathcal{O}_1\}$
- $Fx_{\mathcal{I}}^P(\mathcal{O}_2) = \{O : conj_E(O) \in \mathcal{I} \text{ or (for all clauses } O :- L_1, \dots, L_m \in P \text{ there exists } i, 1 \leq i \leq m, conj_D(L_i) \in \mathcal{I} \text{ or } L_i \in \mathcal{O}_2)\}$

□

The only addition required for explicit negation beyond similar operators for normal programs is the check in the operator $Fx_{\mathcal{I}}^P$ that $conj_E(O) \in \mathcal{I}$, which is used to ensure coherency. Both $Tx_{\mathcal{I}}^P$ and $Fx_{\mathcal{I}}^P$ can be shown to be monotonic and continuous over the information ordering by the usual methods (cf. [28]), leading to the following operator.

Definition 2.2 Let P be a ground program, then ω_{ext}^P is an operator that assigns to every interpretation \mathcal{I}^1 of P a new interpretation \mathcal{I}^2 such that

$$\begin{aligned}\mathcal{I}_T^2 &= lfp(Tx_{\mathcal{I}^1}^P(\emptyset)) \\ \mathcal{I}_F^2 &= \{not(O) | O \in gfp(Fx_{\mathcal{I}^1}^P(objective_literals(P)))\}\end{aligned}$$

□

This latter operation can also be shown to be monotonic over the information ordering of interpretations by the usual methods, leading to the formulation of the well-founded semantics as used in this paper.

Definition 2.3 [Well-founded Semantics for Extended Programs] Let P be a ground extended program. $WFS(P)$ is defined as the least fixed point, over the information ordering, of ω_{ext}^P . □

Example 2.1 Let P be the program containing the rules $\{c :- not(b); \quad b :- a; \quad -b; \quad a :- not(a)\}$. Then $WFS(P) = \{-b, c, not(-a), not(b), not(-c)\}$. Note that to compute c , coherency must be used to infer $not(b)$ from $-b$.

It is important to note that the “model” obtained using ω_{ext}^P may be paraconsistent. Using the operator ω_{ext}^P it is possible to define a stability operator for extended programs that allows partial, and possibly paraconsistent models.

Definition 2.4 [Partial Stable Interpretation of an Extended Program] Let P be a ground extended program. We call an interpretation \mathcal{I} a *partial stable interpretation* of P if $\mathcal{I} = \omega_{ext}^P(\mathcal{I})$ □

If an interpretation \mathcal{I} contains both O and $-O$, then through coherency, $\omega_{ext}^P(\mathcal{I})$ will contain both O and $not(O)$ and so will be inconsistent. Thus, by definition an interpretation \mathcal{I} can be a partial stable interpretation even if it is inconsistent. However as we will see, within abductive frameworks consistency can be ensured by means of integrity constraints — for instance, prohibiting O and $-O$ to be true for any objective literals O . We use $WFS(P)$ as a basis for abduction in part because the support of a literal on a contradiction can be detected by simply looking at the paraconsistent well-founded model. As shown in Sections 6 and 8 of [8] it is the only one of an array of semantics for extended programs with this property, along with having other desirable structural properties.

2.3 Three-Valued Abductive Frameworks

The definitions of three-valued abductive frameworks modify those of [7].

Definition 2.5 [Integrity Rule] An *integrity rule* for a ground program P has the form

$$\perp \text{ :- } L_1, \dots, L_n$$

where each L_i , $1 \leq i \leq n$ is a literal formed over an element of \mathcal{L}_P . \square

Definition 2.6 [Abductive Framework and Abductive Subgoal] An abductive framework is a triple $\langle P, \mathcal{A}, I \rangle$ where \mathcal{A} is a finite set of ground objective literals of \mathcal{L}_P called *abducibles*, such that for any objective literal O , $O \in \mathcal{A}$ iff $\text{conj}_E(O) \in \mathcal{A}$, I is a set of ground integrity rules, and P is a ground program such that (1) there is no rule in P whose head is in \mathcal{A} ; and (2) $\perp / 0$ is a predicate symbol not occurring in \mathcal{L}_P .

An *abductive subgoal* $S = \langle L, \text{Set} \rangle$ is a literal L together with a finite set of abducibles, Set , called the *context* of S . If the context contains both an objective literal and its explicit conjugate, it is termed *inconsistent* and is *consistent* otherwise. \square

Definition 2.6 requires that if an objective literal, say, $-O_1$ is abducible, then O_1 must be as well. This requirement will be used to allow abduction of positive and negative information in a symmetric manner. An abductive subgoal $\langle L, \text{Set} \rangle$ contains a set, Set , of such abducibles, along with a subgoal, L which in the dual programs used by ABDUAL can be a literal. This notation is used to capture the fact that a solution to L is sought in the context in which the (positive and negative) objective literals in Set have been abducted to be true. If $\text{not}(A)$ is a negative literal, occurring in P or I , and A is an abducible objective literal, ABDUAL will provide coherency axioms to propagate the truth value of $-A$ or A to $\text{not}(A)$ if necessary. Thus it is sufficient for the set of abducibles to contain only objective literals. The requirement that there can be no rule in P whose head is an abducible leads to no loss of generality, since any program with abducibles can be rewritten to obey it ⁴.

Definition 2.7 [Abductive Scenario] A scenario of an abductive framework $\langle P, \mathcal{A}, I \rangle$ is a tuple $\langle P, \mathcal{A}, \mathcal{B}, I \rangle$, where $\mathcal{B} \subseteq \mathcal{A}$ is such that there is no $O \in \mathcal{B}$, such that $\text{conj}_E(O) \in \mathcal{B}$. $P_{\mathcal{B}}$ is defined as the smallest set of rules that contains for each $A \in \mathcal{A}$, the rule $A \text{ :- } \mathbf{t}$ iff $A \in \mathcal{B}$; and $A \text{ :- } \mathbf{u}$ otherwise. \square

Definition 2.8 [Abductive Solution] An abductive solution is a scenario $\sigma = \langle P, \mathcal{A}, \mathcal{B}, I \rangle$ of an abductive framework, such that \perp is false in $M(\sigma) = WFS(P \cup P_{\mathcal{B}} \cup I)$. \square

We say that $\sigma = \langle P, \mathcal{A}, \mathcal{B}, I \rangle$ is an abductive solution for a query Q if $M(\sigma) \models Q$. σ is minimal, if there is no other abductive solution $\sigma' = \langle P, \mathcal{A}, \mathcal{B}', I \rangle$ for Q such that $WFS(\mathcal{B}') \subseteq_{\text{info}} WFS(\mathcal{B})$.

The definition of an abductive solution is three-valued in that (objective) literals in P , \mathcal{A} , and I may be undefined. Given a query and an abductive framework, our goal is to construct a solution σ such that

$$M(\sigma) \models Q$$

⁴For instance, if it is desired to make abducible some objective literal A such that A is the head of a rule, one may introduce a new abducible predicate A' , along with a rule $A \text{ :- } A'$. See e.g. [21].

and

$$M(\sigma) \models \text{not } \perp$$

In addition, it is desirable to evaluate only those portions of P and I that are relevant to Q and to construct solutions that are minimal in the sense that as few literals as possible are assigned a value of true or false. Theorem 3.2 below ensures this minimality condition.

3 Query Evaluation over Abductive Solutions

We informally introduce ABDUAL through a series of examples (Formal Definitions can be found in Sections 3.1 and 3.2). ABDUAL shares similarities with SLG in its propagation of delay literals through ANSWER CLAUSE RESOLUTION, in the semantics it attaches to unconditional answers, and in its simplification of delay literals. The first example illustrates these characteristics.

Example 3.1 We first illustrate how ABDUAL can be used to compute queries to ground programs according to the well-founded semantics when neither abduction nor integrity constraints are needed. Accordingly, consider the abductive framework $\langle P_1, \emptyset, \emptyset \rangle$, in which the set of abducibles and the set of integrity rules are both empty, and P_1 is

$$\begin{aligned} p &:- \text{not}(q). \\ p &:- \text{not}(r). \\ q &:- \text{not}(p). \end{aligned}$$

$WFS(P_1)$ restricted to the objective literals $\{p, q, r\}$ is $\{p, \text{not}(q), \text{not}(r)\}$. In order to evaluate the query $?- q$ through ABDUAL, we first create the dual form of P_1 taken together with a *query rule*

$$\text{query} \quad :- \quad q, \text{not}(\perp).$$

where the atom **query** is assumed not to be in \mathcal{L}_{P_1} . This rule ensures that integrity constraints are checked for any abductive solutions that are derived. This dual program, $\text{dual}(\{P_1 \cup \text{query} \quad :- \quad q, \text{not}(\perp)\}, \emptyset)$ is shown in Figure 1.

$\begin{aligned} p &:- \text{not}(q). \\ p &:- \text{not}(r). \\ q &:- \text{not}(p). \\ \\ \text{query} &:- q, \text{not}(\perp). \end{aligned}$	$\begin{aligned} \text{not}(p) &:- q, r. \\ \\ \text{not}(q) &:- p. \\ \text{not}(r) &. \\ \text{not}(\text{query}) &:- \text{not}(q). \\ \text{not}(\text{query}) &:- \perp. \\ \text{not}(\perp) &. \end{aligned}$
$\begin{aligned} \text{not}(p) &:- \neg p. & \text{not}(\neg p) &:- p. \\ \text{not}(q) &:- \neg q. & \text{not}(\neg q) &:- q. \\ \text{not}(r) &:- \neg r. & \text{not}(\neg r) &:- r. \end{aligned}$	

Figure 1: *Dual Program for $P_1 \cup \{\text{query} \quad :- \quad q, \text{not}(\perp)\}$*

Note that in the dual form of a program, P , a rule can have a default literal of the form $\text{not}(A)$ as its head; rules for $\text{not}(A)$ are designed to derive $\text{not}(A)$ if and only if A is false in $WFS(P)$. The last three lines of Figure 1 are *coherency axioms* so-named because they ensure coherency of the model computed by ABDUAL. As is usual with tabled evaluations (e.g. [5]), the ABDUAL evaluation of a query to the above dual program is represented as a sequence $\mathcal{F}_0, \dots, \mathcal{F}_i$, of forests of ABDUAL trees. \mathcal{F}_0 is the forest consisting of the single tree $\langle \text{query}, \emptyset \rangle \text{ :- } | \text{query}$, which sets up resolution for the query rule. Given a successor ordinal $i + 1$, a forest \mathcal{F}_{i+1} is created when an ABDUAL operation either adds a new tree to \mathcal{F}_i or expands a node in an existing tree in \mathcal{F}_i . A forest of trees at the end of one possible ABDUAL evaluation of the above query is shown in Figure 2⁵. Nodes in Figure 2 are all *regular* having the form $\text{Abductive_subgoal} \text{ :- } \text{DelayList} | \text{GoalList}$, where Abductive_subgoal is an abductive subgoal (Definition 2.6), and GoalList and DelayList are both sequences of literals. Intuitively the truth of literals in these sequences must be determined in order to prove or fail the abductive subgoal. When an ABDUAL evaluation encounters a new literal, S , a tree with root $\langle S, \emptyset \rangle \text{ :- } | S$ is added to the forest via the NEW SUBGOAL operation. Thus, in Figure 2, when the literal q is selected in node 1, a NEW SUBGOAL operation creates node 2 as a single tree — indeed, all root nodes other than the initial node 0 are created through one or another application of this operation. Immediate children of the roots of trees are created via PROGRAM CLAUSE RESOLUTION operations, while children of other nodes can be created by a variety of operations to which we now turn.




Figure 2: Simplified ABDUAL Evaluation of a query to $\langle P_1, \emptyset, \emptyset \rangle$

Consider the state of the evaluation after node 5 has been created. The evaluation of q depends on $\text{not}(p)$ and vice-versa. In order to determine the truth of q and $\text{not}(p)$ the literal r must be selected and failed, but this is not possible in a fixed left-to-right selection strategy. The ABDUAL DELAYING operation allows the fixed selection strategy to be broken by moving a selected negative literal from the *GoalList* of a node to its *DelayList* so that further literals in the node, such as r , may be selected. Applied to node 3, the DELAYING operation produces node 6, $\langle q, \emptyset \rangle \text{ :- } \text{not}(p) |$. An *answer* is a regular leaf node with an empty *GoalList*. In the subforest of Figure 2 consisting of nodes whose index is 6 or less, node 6 is an answer. Because its *DelayList* is non-empty it is termed a *conditional answer*. While node 6 is an answer for q , it is not known at the time node 6 is created whether q is true or false — its truth value is conditional on that of $\text{not}(p)$. Answers are returned to other

⁵For simplicity of presentation, Figure 2 does not display computation paths that include the coherency axioms, as they are irrelevant in this example.

nodes via the ANSWER CLAUSE RESOLUTION operation which also combines the abductive contexts of the answer and the node to ensure consistency. Using this operation, the conditional answer is resolved against the selected literal of node 5 producing node 7, $\langle \text{not}(p), \emptyset \rangle :- q|r$. Similarly to SLG, the ANSWER CLAUSE RESOLUTION operation of ABDUAL does not propagate *DelayLists* of conditional answers, thus the literal added to the *DelayList* in node 7 is q rather than the literal originally delayed, $\text{not}(p)$. This action is necessary for ABDUAL to have polynomial complexity for normal programs in the absence of abduction (cf. Theorem 3.4). If *DelayLists* of conditional answers were propagated directly, the number of answers for a given subgoal could be proportional to the number of its derivations (see [5] for an example of such a program). Thus a literal L can be added to the *DelayList* of a node in one of two ways: if L is negative, it can be added through an explicit DELAYING operation; otherwise, L can be added to a *DelayList* if an ANSWER CLAUSE RESOLUTION operation resolves a conditional answer against L regardless of whether L is positive or negative.

Note that after the production of node 8, the evaluation of $\text{not}(p)$ and of all the selected subgoals in the goal list upon which it depends cannot proceed further, and these subgoals cannot produce any new answers, conditional or otherwise. Such subgoals are termed *completely evaluated* (Definition 3.3). At this stage, node 6 contains in its *DelayList* an atom that is known to be false – i.e. that is completely evaluated and has no answers. A SIMPLIFICATION operation is applicable to node 6, creating the failure node, node 8, as its child, so that node 6 is no longer a leaf and hence no longer an answer. After the production of node 9, neither the tree for q nor that for $\text{not}(p)$ has an answer at the end of the evaluation, corresponding to the fact that both literals are false in $WFS(P_1)$.

We now formalize the definitions of some concepts introduced in Example 3.1. For an objective literal O in a program P , $\text{not}(O)$ is defined so that it will be derivable as true iff O is false in $WFS(P)$. For instance, if there is a fact in P for some objective literal O then the dual has no rule for $\text{not}(O)$. The definition below is somewhat more complicated than the form implicitly used in Example 3.1, but as explained below, it ensures both that ABDUAL will be definable on infinite programs and that it will have an appropriate complexity for finite programs.

Definition 3.1 [Dual Program] Let P be a ground extended program, and \mathcal{A} a (possibly empty) finite set of abducibles. The *dual transformation* creates a *dual program* $\text{dual}(P, \mathcal{A})$, defined as the union of P with smallest program containing the sets of rules fold_P and cohere_P as follows:

1. Let O be an objective literal for which there are no facts in P , and with $\beta \leq \omega$ rules of the form:

$$r_i : O :- L_{i,1}, \dots, L_{i,n_i}$$

for $i < \omega$, where each n_i is finite.

- (a) Then fold_P contains the rule

$$\text{not}(O) :- \text{not}(\text{fold}_1^a \text{--} O)$$

along with rules

$$\text{not}(\text{fold}_i^a \text{--} O) :- \text{not}(\text{fold}_i^b \text{--} O), \text{not}(\text{fold}_{i+1}^a \text{--} O)$$

for all i , $1 \leq i < \beta$; and

$$\text{not}(\text{fold}_\beta^a O) :- \text{not}(\text{fold}_\beta^b O)$$

if β is finite.

- (b) and for $1 \leq j \leq \text{maximum}\{n_1, \dots, n_i\}$, such that $L_{i,j}$ exists as a literal in r_i , fold_P contains a rule:

$$\text{not}(\text{fold}_i^b O) :- \text{conj}_D(L_{i,j})$$

where $\text{fold}_k^a O$, $\text{fold}_k^b O$ are assumed not to occur in \mathcal{L}_P for any k (such rules are termed *folding rules*, and literals formed from objective literals whose predicate symbol is $\text{fold}_k^a O$ or $\text{fold}_k^b O$ are called *folding literals*).

2. Otherwise, if $\text{not}(O)$ is in $\text{literals}(P)$, but there is no rule with head O in P , then fold_P contains the rule $\text{not}(O) :- \mathbf{t}$. If there is a fact for O in P , the rule $\text{not}(O) :- \mathbf{f}$ may be introduced or omitted.

3. cohere_P consists of *axioms of coherence* that relate explicit and default negation, defined as:

$$\text{not}(O) :- \text{conj}_E(O)$$

For each objective literal $\text{not}(O)$ in either $\text{literals}(P \cup \text{fold}_P)$ or \mathcal{A} .

□

Example 3.2 Consider a program fragment in which an objective literal \mathbf{m} is defined as:

$$\begin{aligned} \mathbf{m} & :- \mathbf{n}_1, \text{not}(\mathbf{o}_1). \\ \mathbf{m} & :- \mathbf{n}_2, \text{not}(\mathbf{o}_2). \\ \mathbf{m} & :- \mathbf{n}_3, \text{not}(\mathbf{o}_3). \end{aligned}$$

Note that a naive dualization of \mathbf{m} as implicitly used in Example 3.1 (and as defined in Definition A.1) would produce a rule for each partial truth assignment to the body literals of \mathbf{m} that falsifies \mathbf{m} , leading to 8 rules, each with 3 body literals. Indeed, it is easy to see that naive dualization of a predicate p with β clauses can lead to a predicate for $\text{not}(p)$ that has a number of clauses exponential in β , making the naive dual form unsuitable in terms of complexity for finite programs. Furthermore, the number of body literals in a clause for $\text{not}(p)$ may be linear in β so that if the naive transformation were used, the dual of an infinite program would not be a program as defined in Section 2.1.

The folding rules in the dual form of m (fold_P of Definition 3.1) are shown in Figure 3. In Definition 3.1, if there are an infinite number of rules defining an objective literal O , there will also be an infinite number of folding rules defining $\text{not}(O)$, but each rule will have a finite sequence of literals in their body. Also note that in a finite ground program, if an objective literal O is defined by n rules each of which have m body literals, the size of the rules defining O will be $mn + n$ (see Definition A.4 for a precise definition of the size of rules and programs). In $\text{dual}(P, \mathcal{A})$, there will be $m * n$ rules of the form $\text{fold}_i^b O_i$ for some i , each of size 2, along with folding rules of the type $\text{fold}_i^a O_i$ for some i so that the size of the rules for $\text{not}(O)$ in $\text{dual}(P, \mathcal{A})$ is linear in the size of the rules for O in P .

```

not(m)    :- not(fold1am).

not(fold1am) :- not(fold1bm), not(fold2am).
not(fold2am) :- not(fold2bm), not(fold3am).
not(fold3am) :- not(fold3bm).

not(fold1bm) :- not(n1).
not(fold1bm) :- o1.
not(fold2bm) :- not(n2).
not(fold2bm) :- o2.
not(fold3bm) :- not(n3).
not(fold3bm) :- o3.

```

Figure 3: *Folded Dual Program for a program clause*

While the dual form of Definition 3.1 is necessary for the correctness and complexity results that follow, examples will use a simpler form without folding literals that is logically equivalent for finite programs (see Definition A.1 for an exact statement of this simpler form).

Definition 3.2 [ABDUAL Trees and Forest] An ABDUAL forest consists of a forest of ABDUAL trees. Nodes of ABDUAL trees are either *failure nodes* of the form *fail*, or *regular nodes* of the form

$$Abductive_Subgoal \text{ :- } DelayList | GoalList$$

where *Abductive_subgoal* is an abductive subgoal (Definition 2.6). Both *DelayList* and *GoalList* are finite sequences of literals (also called *delay literals* and *goal literals*, respectively).

We call a regular leaf node *N* an *answer* when *GoalList* is empty. If *DelayList* is also empty, *N* is *unconditional*; otherwise it is *conditional*. \square

Definition 3.9 will ensure that the root node of a given ABDUAL tree, *T*, has the form $\langle S, \emptyset \rangle \text{ :- } |S$, where *S* is a literal. In this case, we say that *S* is the *root goal* for *T* or that *T* is the *tree for S*. Similarly by Definition 3.9, a forest contains a root goal *S* if the forest contains a tree for *S*. Literal selection rules apply to the *GoalList* of a node; as mentioned in Section 2, we use a fixed left-to-right order for simplicity of presentation so that the leftmost literal in the *GoalList* of a node is termed the *selected literal* of the node.

Example 3.3 The well-founded semantics captures infinite recursion by means of the concept of unfounded sets: an atom involved in an unfounded set is assigned a truth-value of false. When a program undergoes the dual transformation, negative literals involved in infinite recursion must be made to succeed. As an example of this, consider the abductive framework $\langle P_2, \emptyset, \emptyset \rangle$ in which *P*₂ is defined as:

```

s  :- not(p), not(q), not(r).
p  :- not(s), not(r), q.
q  :- not(p), r.
r  :- not(q), p.

```

Note that $WFS(P_2)$ restricted to the objective literals $\{s, p, q, r\}$ is $\{s, \text{not}(p), \text{not}(q), \text{not}(r)\}$. Assuming the query $?- s$ to $\langle P_2, \emptyset, \emptyset \rangle$, the dual program $\text{dual}((\{P_2 \cup \text{query} :- s, \text{not}(\perp)\}), \emptyset)$ is shown in Figure 4 ⁶

<pre> s :- not(p), not(q), not(r). p :- not(s), not(r), q. q :- not(p), r. r :- not(q), p. query :- s, not(\perp). </pre>	<pre> not(s) :- p. not(s) :- q. not(s) :- r. not(p) :- s. not(p) :- r. not(p) :- not(q). not(q) :- p. not(q) :- not(r). not(r) :- q. not(r) :- not(p). not(query) :- not(s). not(query) :- \perp. not(\perp). </pre>
--	--

<pre> not(p) :- -p. not(q) :- -q. not(r) :- -r. not(s) :- -s. </pre>	<pre> not(-p) :- p. not(-q) :- q. not(-r) :- r. not(-s) :- s. </pre>
--	--

Figure 4: *Dual Program for $P_2 \cup \{\text{query} :- s, \text{not}(\perp)\}$*

An ABDUAL forest at the end of an evaluation of *query* is shown in Figure 5. As can be seen from Figure 5, the evaluation at first proceeds using the same operations as in Example 3.1, where the roots of non-initial trees are created via NEW SUBGOAL operations, the children of roots of trees created via PROGRAM CLAUSE RESOLUTION operations, and other nodes created via ANSWER CLAUSE RESOLUTION or DELAYING operations. However node 27 is produced by a new operation. Note that in the subforest of Figure 5 consisting of nodes numbered 26 or less, nodes 24, 25, and 26 are all conditional answers that “depend” on each other through their *DelayLists*. However, in the well-founded model of P_2 , p , q and r should be false as they belong to an unfounded set (based on the empty interpretation). In order to derive their truth-values ABDUAL includes a CO-UNFOUNDED SET REMOVAL operation. Nodes 24, 25, and 26 together form an analogue in the dual program to an unfounded set [35] consisting of p , q , and r in P_2 . Such an analogue is called a *co-unfounded set*. Whereas positive literals in an unfounded set are all false, negative literals in a co-unfounded set are

⁶The transformation in Definition A.1 is used for simplicity.

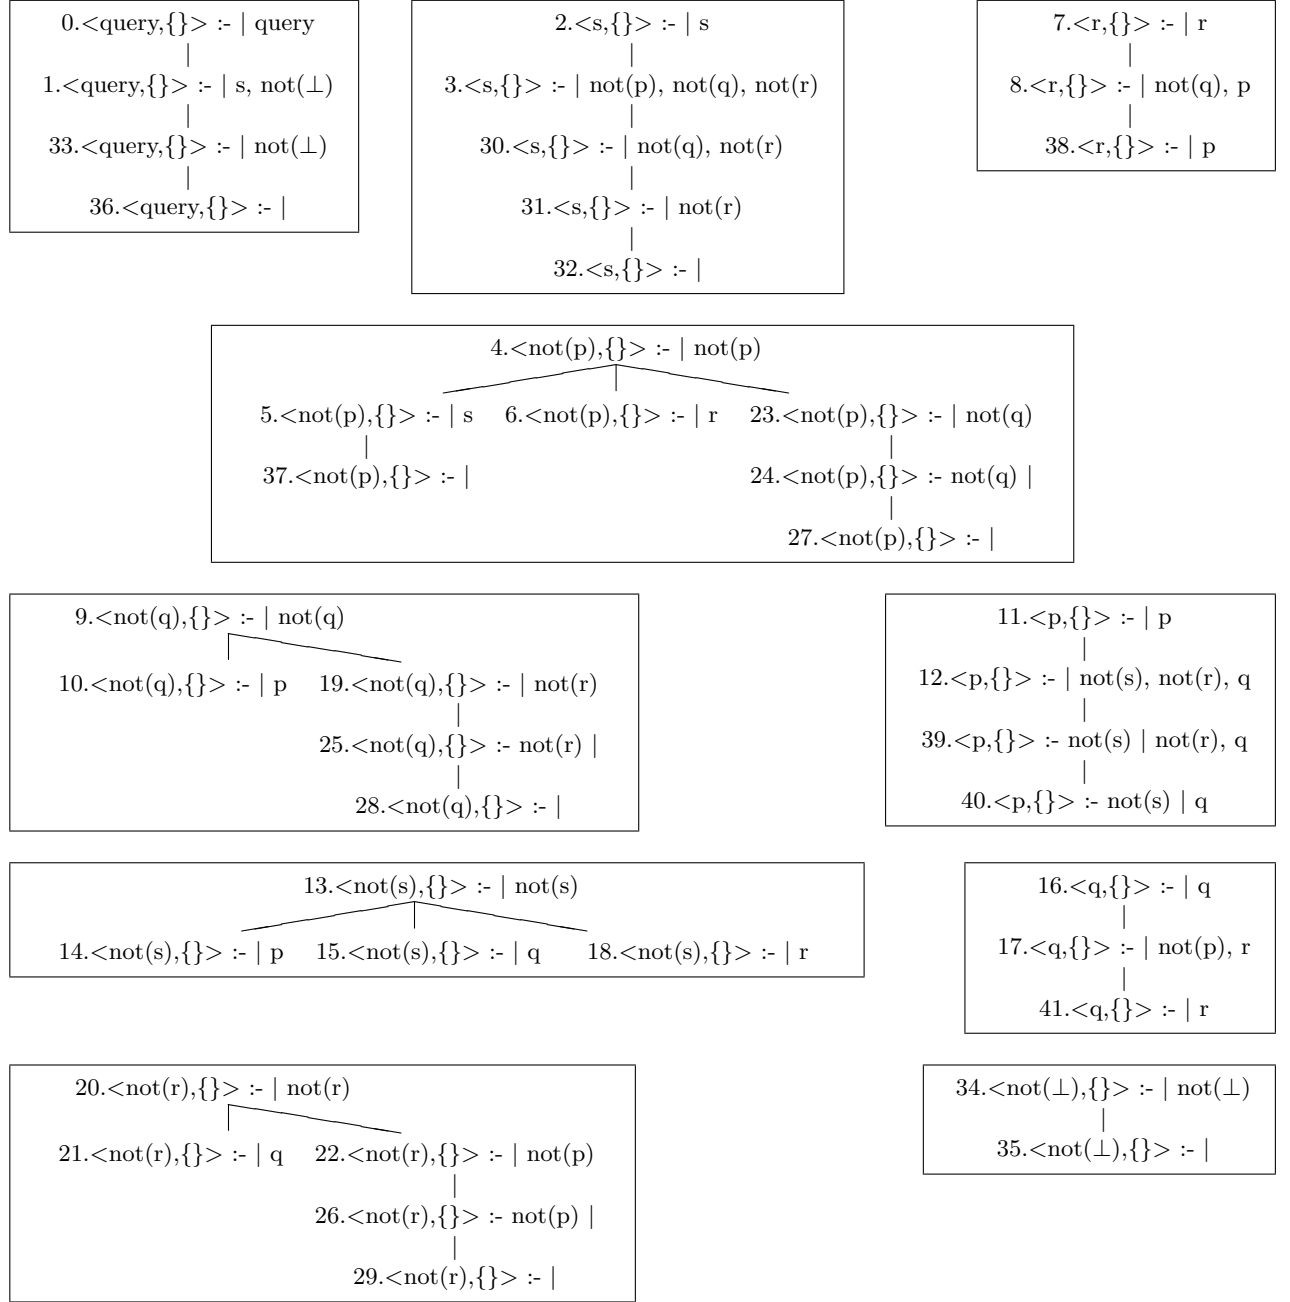


Figure 5: Simplified ABDUAL Evaluation of a query to $\langle P_2, \emptyset, \emptyset \rangle$

all true. When an answer is determined to belong to a co-unfounded set, it is made unconditionally true. In this example the CO-UNFOUNDED SET REMOVAL operation creates the unconditional answer, node 27, while SIMPLIFICATION operations produce nodes 28 and 29. Subsequently, ANSWER CLAUSE RESOLUTION resolves the (unconditional) answer $\langle \text{not}(p), \emptyset \rangle$ against the selected literal of node 3 to create node 30 through ANSWER CLAUSE RESOLUTION, and subsequent applications of this operation produce nodes 31-34, 36-38, 40 and 41.

We summarize some of the elements of the previous two examples. Intuitively, the distinction between goal literals and delay literals is that goal literals are currently selected within a node or are yet to be selected. As a result, there is an answer for S if there is a regular leaf node N in a tree for S that has no goal literals. If N does not contain delay literals, it is an unconditional answer and S has an abductive solution defined by the context of the abductive subgoal of N ; if N does contain delay literals, then it is a conditional answer and the abductive context for N is not yet determined by the evaluation to make N either true or false. Finally, a completely evaluated subgoal (Definition 3.3) that has no answers at all is interpreted to be false for all abductive contexts. At an operational level, as described in Definition 3.9, goal literals may be resolved away via an ANSWER CLAUSE RESOLUTION operation, abduced, or delayed. Literals in the *DelayList* were not resolved away when they were selected, but rather their resolution was postponed. Delay literals are subject to the CO-UNFOUNDED SET REMOVAL operation mentioned in Example 3.3, and also to SIMPLIFICATION operations mentioned in Example 3.1. Maintaining both delay literals and goal literals within an evaluation is useful as it is necessary to identify unfounded sets of objective literals within the well-founded semantics, as well as co-unfounded sets of objective literals within the dual form of a program. Determining (co-)unfounded sets is expensive in practical terms, so that restricting such an operation to delay literals can form an important optimization (cf. [11]).

The notion of a set of ABDUAL trees being completely evaluated was introduced in Example 3.1 to capture the concept of when a set of trees in a forest has returned all of the answers in the model of a program. This can happen in one of two ways. First, a tree may contain an unconditional answer whose abductive context is empty, in which case further evaluation will not produce any more minimal abductive answers. Second, a tree may have had all possible ABDUAL operations performed on the selected literal in the *GoalList* of each of its nodes. For this condition to occur, all possible answers must have been returned to the selected literals so that a tree is not completely evaluated unless all trees that it depends on (through the selected literal of each of its nodes) are completely evaluated as well. An example of this occurs in Example 3.3, where the trees for p , q , and r are mutually dependent and may only be evaluated together.

Definition 3.3 [Completely Evaluated] Given an ABDUAL forest \mathcal{F} , a set \mathcal{T} of ABDUAL trees is *completely evaluated* iff at least one of the following conditions is satisfied for each tree $T \in \mathcal{T}$:

1. T contains an unconditional answer whose abductive subgoal context is empty; or
2. For each node N in \mathcal{T} with selected goal literal SL
 - The tree for SL belongs to a set S' of completely evaluated trees; and

- No NEW SUBGOAL, PROGRAM CLAUSE RESOLUTION, ANSWER CLAUSE RESOLUTION, DELAYING, or ABDUCTION operations (Definition 3.9) are applicable to N .

A literal L is completely evaluated in \mathcal{F} if the tree for L belongs to a completely evaluated set in \mathcal{F} .
 \square

Finally, we turn to an example to illustrate how ABDUAL can evaluate queries to general abductive frameworks.

Example 3.4 Consider the abductive framework $\langle P_3, \mathcal{A}_3, I_3 \rangle$, in which P_3 is the program

```
p :- not(q*).
q :- not(p*).
```

$\mathcal{A}_3 = \{p^*, q^*, -p^*, -q^*\}$, and I_3 is the program

```
⊥ :- p_constr
⊥ :- q_constr
p_constr :- p, -p*.
q_constr :- q, -q*.
```

So that the (ground) integrity constraints represent an abductive interpretation of default negation. Let the query rule be

```
query :- q, not(⊥).
```

The dual program with coherency axioms (simplified for presentation by using the transformation of Definition A.1, which does not include folding predicates.) is shown in Figure 6.

p :- not(q*).	not(p) :- q*.
q :- not(p*).	not(q) :- p*.
⊥ :- p_constr	not(⊥) :- not(p_constr), not(q_constr)
⊥ :- q_constr	
p_constr :- p, -p*.	not(p_constr) :- not(p)
	not(p_constr) :- not(-p*).
q_constr :- q, -q*.	not(q_constr) :- not(q)
	not(q_constr) :- not(-q*).
query :- q, not(⊥).	not(query) :- not(q).
	not(query) :- ⊥.
not(-p) :- p.	not(p) :- -p.
not(-q) :- q	not(q) :- -q
not(-p*) :- p*	not(p*) :- -p*
not(-q*) :- q*	not(q*) :- -q*
not(-p_constr) :- p_constr	not(p_constr) :- -p_constr
not(-q_constr) :- q_constr	not(q_constr) :- -q_constr

Figure 6: *Dual program for $P_3 \cup \{query :- q, not(\perp)\} \cup \mathcal{A} \cup I$*

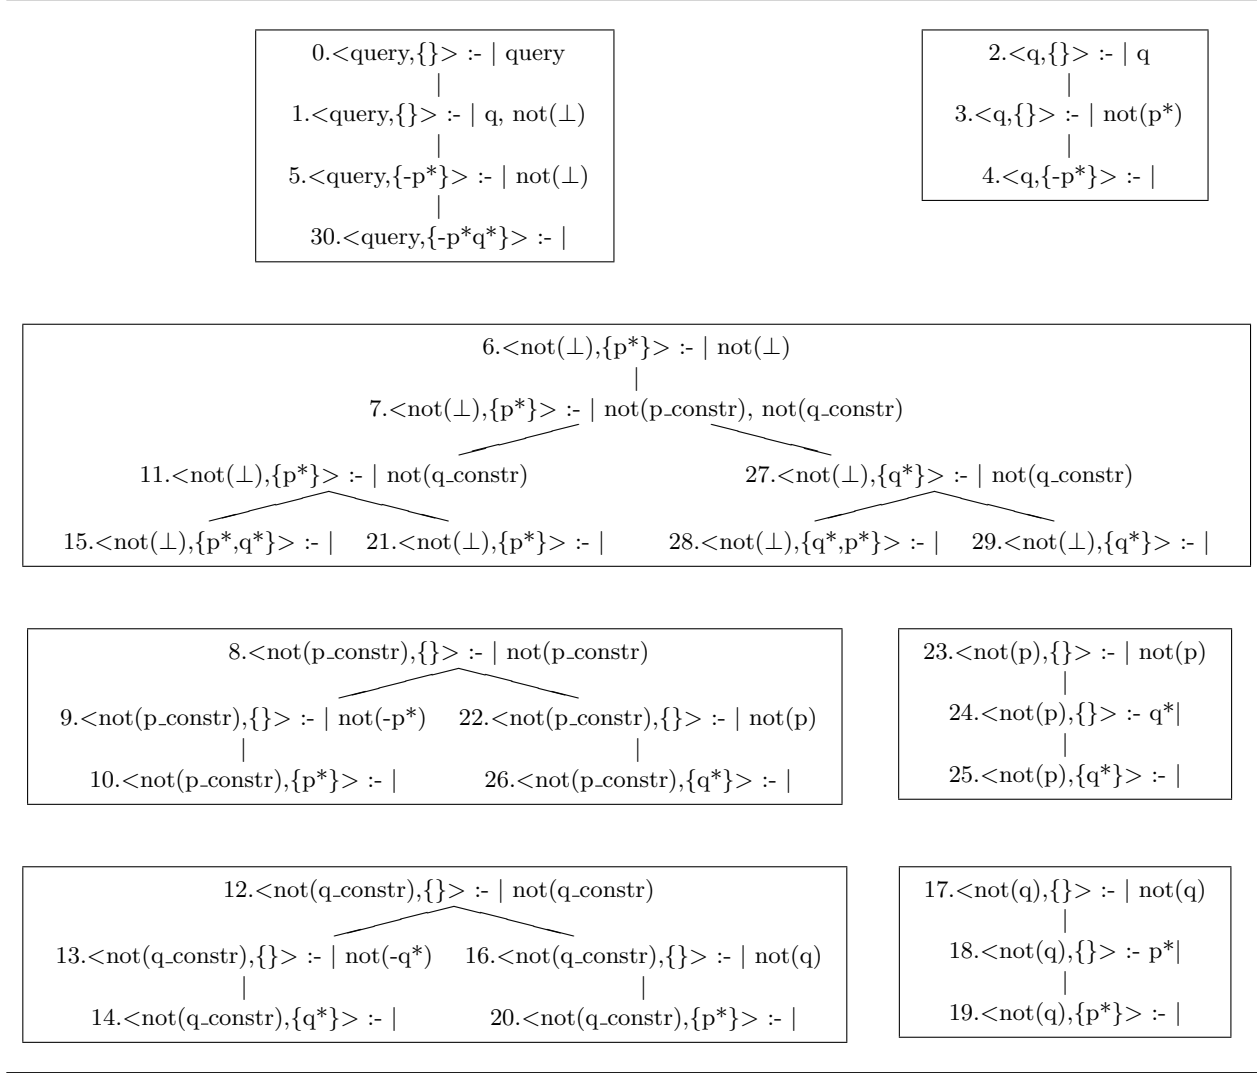


Figure 7: *Simplified ABDUAL evaluation of a query to $\langle P_3, \mathcal{A}_3, I_3 \rangle$.*

Figure 7 illustrates a forest of trees created by an ABDUAL evaluation of this initial query. For purposes of space, it does not depict derivations stemming from coherency axioms. When an abductive framework contains a non-trivial set of abducibles, provision must be made for when the selected literal of a given node is an abducible, as well as for propagating abducibles among abductive subgoals. In the first case, if the selected literal of a node N is an abducible, and the addition of the selected literal to the context of the abductive subgoal of N does not make the context inconsistent (Definition 2.6), an ABDUCTION operation is applicable to N . For instance, ABDUCTION operations are used to produce nodes 25 and 19. The figure also illustrates cases in which abducibles are propagated through ANSWER CLAUSE RESOLUTION. Node 5 is produced by resolving the answer $\langle q, \{-p^*\} \rangle \text{ :- } \mid$ against the selected literal, q of node 1, to produce the (consistent) context $\{-p^*\}$. Abducibles therefore differ from delay literals in that the abducibles are propagated into the context of an abductive subgoal, while delayed literals are not propagated into delay lists. Propagating abducibles through ANSWER CLAUSE RESOLUTION operations is common in this derivation, producing in a similar manner, nodes 20, 4, 11, 27, 15, 21, 28, 29, 9, 26, and 14.

Certain of these nodes are created using the coherency axioms, which are not shown in Figure 7. For instance in producing node 4, $\langle q, \{-p^*\} \rangle \text{ :- } \mid$, a NEW SUBGOAL operation creates a new tree for the selected literal, $\text{not}(p^*)$ of node 3. This tree uses the rule $\text{not}(p^*) \text{ :- } \neg p^*$ for PROGRAM CLAUSE RESOLUTION, and then abduces $\neg p^*$, propagating the abducible to the context of node 4. In propagating abducibles, the ANSWER CLAUSE RESOLUTION operation enforces the restriction that the context of the answer must be consistent with the context of the abductive subgoal of the node to which the answer is returned. For instance, of the two unique abductive solutions to $\text{not}(\perp)$ only one can be returned to the node $\langle \text{query}, \{-p^*\} \rangle \text{ :- } \mid \text{not}(\perp)$, namely $\langle \text{not}(\perp), \{q^*\} \rangle$.

The final definitions for ABDUAL are now provided, beginning with the unfounded and co-unfounded sets.

3.1 Unfounded and Co-unfounded Sets

One of the ideas behind of the well-founded semantics of normal programs is to assign the value of false to atoms that are contained in unfounded sets. Intuitively these sets can be seen as including atoms whose derivations lead to positive loops or to infinite chains of dependencies among subgoals. Unfounded sets for extended logic programs are defined as follows:

Definition 3.4 [Unfounded Set of Objective Literals] Let P be a ground extended logic program, and \mathcal{I} a coherent interpretation of P . Then a set of objective literals $\mathcal{S} \subseteq \text{literals}(P)$ is an unfounded set of P with respect to \mathcal{I} if for each rule r_s with head $H \in \mathcal{S}$ one of the following conditions hold:

1. for some body literal L_i in r_s , the default conjugate of L_i is in \mathcal{I} .
2. for some positive body literal L_i in r_s , $L_i \in \mathcal{S}$.

A literal that makes either condition true is called a *witness of unusability* for rule r_s with respect to \mathcal{I} . □

A witness of unusability in \mathcal{I} may be a literal that is false in \mathcal{I} , or a positive literal whose proof depends on positive literals that are neither contained in \mathcal{I} nor are provable from literals in \mathcal{I} . For instance, in the program $\{p :- q, q :- p\}$ both p and q are unfounded when $\mathcal{I} = \emptyset$. In a dual program, there is also the dual notion of a *co-unfounded set of literals*.

Definition 3.5 [Co-unfounded set of literals] Let P be a ground extended program, \mathcal{A} a set of abducibles, and \mathcal{I} a coherent interpretation of $dual(P, \mathcal{A})$. Then a set of negative literals $\mathcal{S} \subseteq literals(dual((P, \mathcal{A}))$ is a co-unfounded set with respect to \mathcal{I} if for each $H \in \mathcal{S}$, there is a rule $H :- Body$ such that for each $L_i \in Body$:

1. L_i is true in \mathcal{I} ; or
2. $L_i \in \mathcal{S}$.

□

Just as unfounded sets of objective literals are false in a program, co-unfounded sets of negative literals are true in the dual of a program (cf. Lemma A.1). Because any selected negative literal can be delayed, ABDUAL need only take account of co-unfounded sets of literals that occur in *DelayLists* of nodes. For instance, evaluation of the program in Example 3.3 required detection of a co-unfounded set among literals in the *DelayLists* of nodes 24, 25, and 26. A *co-unfounded set of answers* corresponds to a co-unfounded set of literals that arises in a ABDUAL evaluation, and is defined as:

Definition 3.6 [Co-unfounded Set of Answers] Let \mathcal{F} be an ABDUAL forest, and \mathcal{S} a non-empty set of answers in \mathcal{F} . Then \mathcal{S} is a *co-unfounded set* in \mathcal{F} iff

1. Each literal S_i , such that $\langle S_i, C_i \rangle$ is the abductive subgoal of an answer in \mathcal{S} , is a completely evaluated negative literal. Further, S_i is contained in the *DelayList* of some answer in \mathcal{S} .
2. The set

$$Context = \bigcup \{C_i | \langle S_i, C_i \rangle :- DL | \text{ is an answer in } \mathcal{S}\}$$

is consistent; and

3. For each answer $\langle S_i, C_i \rangle :- DL_i | \in \mathcal{S}$

- (a) DL_i is non-empty; and
- (b) for each $S_j \in DL_i$, there exists an answer $\langle S_j, Context_j \rangle :- DL_j | \in \mathcal{S}$.

□

The requirement in condition 1 of Definition 3.6 that the literals be completely evaluated is for convenience, so that an evaluation need not detect co-unfounded sets of answers when more direct derivations may still be possible.

Analogous to a co-unfounded set of answers are the non-supported objective literals. Intuitively, non-supported literals in an ABDUAL forest correspond to unfounded objective literals under a given interpretation.

Definition 3.7 [Supported Objective Literals] Let \mathcal{F} be a forest, and S a positive literal that is the root goal for a tree T in \mathcal{F} . Then S is supported in \mathcal{F} iff

1. T is not completely evaluated; or
2. T contains an answer $\langle S, Context \rangle :- DL$ in T with no positive delay literals in DL ; or
3. T contains an answer $\langle S, Context \rangle :- DL$ in T such that, every positive delay literal L_1 in DL is supported in \mathcal{F} .

□

A tree in a forest is thus supported if it is not completely evaluated, if it contains an unconditional answer, if it contains an answer with a delayed negative literal, or if it contains an answer containing positive literals all of which are themselves supported. The SIMPLIFICATION operation of Definition 3.9 removes an answer of an unfounded literal from a forest by creating a failure node as a child of the answer.

3.2 ABDUAL Evaluations and Operations

An ABDUAL evaluation consists of a (possibly transfinite) sequence of ABDUAL forests ⁷. In order to define the behavior of an ABDUAL evaluation at a limit ordinal, we define a notion of a least upper bound for a set of ABDUAL trees. Any rooted tree can be viewed as a partially ordered set in which each node N is represented as $\{N, P\}$ in which P is a tuple representing the path from N to the root of the tree. When represented in this manner, it is easily seen that when T_1 and T_2 are rooted trees, $T_1 \subseteq T_2$ iff T_1 is a subtree of T_2 , and furthermore, that if T_1 and T_2 have the same root, their union can be defined as their set union, for T_1 and T_2 taken as sets. However, we will sometimes abuse notation in our definitions and refer to trees using the usual graph-theoretic terminology.

Definition 3.8 [ABDUAL Evaluation] Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework and Q a query. An ABDUAL evaluation \mathcal{E} of Q to $\langle P, \mathcal{A}, I \rangle$ is a sequence of ABDUAL forests $\mathcal{F}_0, \mathcal{F}_1, \dots, \mathcal{F}_n$ operating on the ground instantiation of $dual((P \cup I \cup \{query :- not(\perp), Q\}), \mathcal{A})$ such that:

- \mathcal{F}_0 is the forest containing the single tree, $\langle query, \emptyset \rangle :- query$,
- For each successor ordinal $n + 1$, \mathcal{F}_{n+1} is obtained from \mathcal{F}_n by applying an ABDUAL operation from Definition 3.9.
- For each limit ordinal α , \mathcal{F}_α is defined such that $T \in \mathcal{F}_\alpha$ iff
 - The root node of T , $\langle S, \emptyset \rangle :- S$ is the root node of some tree in a forest \mathcal{F}_i , $i < \alpha$;
 - $T = \cup_{i < \alpha} (\{T_i | T_i \in \mathcal{F}_i \text{ and } T_i \text{ has root } \langle S, \emptyset \rangle :- S\})$

If no operation is applicable in \mathcal{F}_n , then it is called a *final forest* of \mathcal{E} .

□

⁷Our definition here follows that of [33] for generalized SLG trees.

In accordance with Definition 3.8, the following ABDUAL operations operate on dual programs.

Definition 3.9 [ABDUAL Operations] Let \mathcal{F}_n be an ABDUAL forest for an evaluation of a query Q to an abductive framework $\langle P, \mathcal{A}, I \rangle$, and suppose $n + 1$ is a successor ordinal. Then \mathcal{F}_{n+1} may be produced by one of the following operations

1. NEW SUBGOAL: Let \mathcal{F}_n contain a non-root node

$$N = \langle S, Context \rangle \text{ :- } DL|L, GoalList.$$

If L is not an abducible and \mathcal{F}_n contains no tree with root goal L , add the tree: $\langle L, \emptyset \rangle \text{ :- } |L$.

2. PROGRAM CLAUSE RESOLUTION: Let \mathcal{F}_n contain a root node

$$N = \langle S, \emptyset \rangle \text{ :- } |S$$

and let there be a clause $S \text{ :- } Body$ in the dual program. If in \mathcal{F}_n , N does not have a child:

$$N_{child} = \langle S, \emptyset \rangle \text{ :- } |Body$$

then add N_{child} as a child of N .

3. ANSWER CLAUSE RESOLUTION: Let \mathcal{F}_n contain a non-root node

$$N = \langle S, Context_1 \rangle \text{ :- } DL_0|L, Body$$

and suppose that \mathcal{F}_n contains an answer node $\langle L, Context_2 \rangle \text{ :- } DL_1|$, such that $Context_1 \cup Context_2$ is consistent. Let $DL_2 = DL_0, L$ if DL_1 is not empty, and $DL_2 = DL_0$ otherwise. Finally, if in \mathcal{F}_n , N does not have a child

$$N_{child} = \langle S, Context_1 \cup Context_2 \rangle \text{ :- } DL_2|Body$$

then add N_{child} as a child of N .

4. DELAYING: Let \mathcal{F}_n contain a non-root leaf node

$$N = \langle S, Context \rangle \text{ :- } DL|not(L), Body$$

where L is not an abducible, and where \mathcal{F}_n contains a tree for $not(L)$, but no answer of the form $\langle not(L), \emptyset \rangle \text{ :- } |$. Then add: $\langle S, Context \rangle \text{ :- } DL, not(L)|Body$ as a child of N .

5. SIMPLIFICATION: Let $N = \langle S, Context_1 \rangle \text{ :- } DL|$ be a node for a tree with root goal S , and let D be a delay literal in DL . Then

- if \mathcal{F}_n contains an unconditional answer node $\langle D, Context_2 \rangle \text{ :- } |$, and if $Context_1 \cup Context_2$ is consistent, let $DL_1 = DL - D$. If

$$N_{child} = \langle S, Context_1 \cup Context_2 \rangle \text{ :- } DL_1|$$

is not a descendant of N in \mathcal{F} , add N_{child} as a child of N .

- if the tree for D is completely evaluated and contains no answers whose context is consistent with C_1 ; or if D is a positive literal that is non-supported, then create a child *fail* of N .

6. CO-UNFOUNDED SET REMOVAL: Let

$$N = \langle S, Context_S \rangle \text{ :- } DL|$$

be an answer in \mathcal{F}_n , such that there is a minimal co-unfounded set of answers \mathcal{S} in \mathcal{F}_n containing N together with answers $\langle L_i, Context_i \rangle \text{ :- } DL_i|$ for all literals $L_i \in DL$. Let

$$Context_{union} = Context_S \cup \bigcup_{\langle L_i, Context_i \rangle \text{ :- } DL_i \in \mathcal{S}} Context_i$$

Then if N does not have a child $N_{child} = \langle S, C_{union} \rangle \text{ :- } |$, create a child N_{child} of N .

7. ABDUCTION: Let

$$N = \langle S, Context \rangle \text{ :- } DL|A, Body$$

where A is an abducible and suppose that $\{A\} \cup Context$ is consistent. Finally, assume that in \mathcal{F}_n , N does not have a child

$$N_{child} = \langle S, Context \cup \{A\} \rangle \text{ :- } DL|Body$$

Then add N_{child} as a child of N .

□

For a discussion of the similarities between definitions of ABDUAL and those of SLG see Section 5.

3.3 Soundness and Completeness of ABDUAL

The first result on the correctness of ABDUAL concerns the correctness of the dual transformation itself (Definition 3.1). To show this, we introduce fixed point operators for dual programs that are analogous to those of Section 2.2, and show that they can be used to construct the well-founded semantics. As before, there are two sets of operators. The first operators form the inner fixed point for dual programs, and are analogous to the operators of Definition 2.1.

Definition 3.10 For a ground program P and set \mathcal{A} of abducibles, let $dual(P, \mathcal{A})$ be the dual program formed by applying the dual transformation to P and \mathcal{A} . Let \mathcal{O}_1 be a set of ground objective literals formed over $dual(P, \mathcal{A})$ and \mathcal{L}_1 be a set of ground negative literals formed over $dual(P, \mathcal{A})$.

- $Td_{\mathcal{I}}^{dual(P, \mathcal{A})}(\mathcal{O}_1) = \{O | O \text{ is an objective literal and there is a clause } O \text{ :- } L_1, \dots, L_n \in dual(P, \mathcal{A}) \text{ and for each } 1 \leq i \leq n, L_i \in \mathcal{I} \text{ or } L_i \in \mathcal{O}_1\}$
- $Fd_{\mathcal{I}}^{dual(P, \mathcal{A})}(\mathcal{L}_1) = \{not(O) | O \text{ is an objective literal and there is a clause } not(O) \text{ :- } L_1, \dots, L_n \in dual(P, \mathcal{A}) \text{ and for each } 1 \leq i \leq n, L_i \in \mathcal{I} \text{ or } L_i \in \mathcal{L}_1\}$

□

The operator $Td_{\mathcal{I}}^{dual(P, \mathcal{A})}$ is essentially the same as the operator $Tx_{\mathcal{I}}^P$ of Definition 2.1; however the operator $Fd_{\mathcal{I}}^{dual(P, \mathcal{A})}$ differs significantly from its analogue, since negative literals are defined by clauses in the dual program that are to be made true. As before, these operators can be shown to be monotonic by the usual methods, leading to the following operator which is used in the outer fixed point.

Definition 3.11 For a ground program P and set \mathcal{A} of abducibles, let $dual(P, \mathcal{A})$ be the dual program formed by applying the dual transformation to P and \mathcal{A} . Let

$$\mathcal{S}_{dual(P, \mathcal{A})} = \{not(O) \mid not(O) \in literals(dual(P, \mathcal{A}))\}$$

$\omega_d^{dual(P, \mathcal{A})}$ is an operator that assigns to every interpretation \mathcal{I}^1 of P a new interpretation \mathcal{I}^2 such that

$$\begin{aligned}\mathcal{I}_T^2 &= lfp(Td_{\mathcal{I}^1}^P(\emptyset)) \\ \mathcal{I}_F^2 &= gfp(Fd_{\mathcal{I}^1}^P(\mathcal{S}_{dual(P, \mathcal{A})}))\end{aligned}$$

□

In the above definition, the use of the greatest fixed point of the operator $Fd_{\mathcal{I}^1}^P(\mathcal{S}_{dual(P, \mathcal{A})})$ captures the fact that co-unfounded sets need to be made true when evaluating the dual program. The following theorem shows the correctness of the dual transformation.

Theorem 3.1 For a ground program P and empty set of abducibles, let $dual(P, \emptyset)$ be the dual program formed by applying the dual transformation to P . Then for any literal $L \in literals(P)$,

$$L \in WFS(P) \iff L \in lfp(\omega_d^{dual(P, \emptyset)})$$

where the least fixed point of $lfp(\omega_d^{dual(P, \emptyset)})$ is taken with regard to the information ordering of interpretations.

Recall that for a query Q to an abductive framework, an ABDUAL evaluation as stated in Definition 3.8 represents solutions to Q that are consistent with the integrity rules by means of a query rule $query \text{ :- } Q, not(\perp)$. The following theorem shows that the ABDUAL operations, acting on the dual program correctly compute abductive solutions.

Theorem 3.2 Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework, and \mathcal{E} be an ABDUAL evaluation of Q against $\langle P, \mathcal{A}, I \rangle$. Then

- \mathcal{E} will have a final forest \mathcal{E}_β ;
- if $\langle query, Set \rangle \text{ :- } \mid$ is an answer in \mathcal{E}_β , $\sigma = \langle P, \mathcal{A}, Set, I \rangle$ is an abductive solution for $\langle P, \mathcal{A}, I \rangle$;
- if $\langle P, \mathcal{A}, Set, I \rangle$ is a minimal abductive solution for Q , then $\langle query, Set \rangle \text{ :- } \mid$ is an answer in \mathcal{E}_β .

Proof: The proof is contained in the Appendix.

◇

3.4 Finite Termination and Complexity of ABDUAL for Extended Programs

Termination of ABDUAL evaluations is guaranteed under the following conditions.

Theorem 3.3 *Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework such that P and I are finite ground extended programs, and \mathcal{A} is a finite set of abducibles. Let \mathcal{E} be an ABDUAL evaluation of a query Q against $\langle P, \mathcal{A}, I \rangle$. Then \mathcal{E} will reach a final forest after a finite number of ABDUAL operations.*

Proof: The proof is contained in the Appendix. \diamond

It is known that the problem of query evaluation to abductive frameworks is NP-complete, even for those frameworks in which entailment is based on the well-founded semantics [14]. More precise results can be obtained for ABDUAL, as shown in the following theorem, which uses a summation over abductive contexts (using of the combinatorial selection function “choose”) to determine the cost of an ABDUAL evaluation. The following theorem relies on a definition of size that is made precise in Definition A.4.

Theorem 3.4 *Let \mathcal{F} be the final forest in an ABDUAL evaluation \mathcal{E} of a query Q against a finite ground abductive framework $\langle P, \mathcal{A}, I \rangle$. Let $C_{context}$ be the maximal cardinality of the context of any abductive subgoal in \mathcal{F} , and $C_{abducibles}$ be the cardinality of \mathcal{A} . Then \mathcal{F} can be constructed in $M \times 2 \times \text{size}(\text{dual}((P \cup I), \mathcal{A}))$ steps, where*

$$M = \sum_{i \leq C_{context}} \binom{C_{abducibles}}{i}$$

Proof: The proof is given in the Appendix. \diamond

Intuitively, this theorem states that the complexity of an ABDUAL evaluation is proportional to the maximal number of abducibles in any abductive subgoals, and to the number of abducibles in the framework. If the number of either of these factors can be reduced, then the complexity of the evaluation will be reduced. Since the size of a dual program is linear in the size of an abductive framework (cf. Lemma A.5), a corollary of Theorem 3.4 is that if the set of abducibles and integrity rules are both empty, the final forest of a ABDUAL evaluation requires a number of operations that is linear in the size of the input program. It is important to note, however, that ABDUAL operations may not be implementable with constant cost. In particular, some operations such as CO-UNFOUNDED SET REMOVAL or removal of a non-supported answers may require a cost that is linear in the size of a program so that the cost of evaluating a program with empty abducibles and integrity rules may not be linear but will remain polynomial (see [11] for an extended discussion of costs of tabling normal programs).

Theorem 3.4 can be used to show that abduction over the well-founded semantics is *fixed-parameter tractable* [12]. Recall that a decision problem Pr can be defined as a question to be answered with a “yes” or “no”. This question has several input parameters. If particular values of these input parameters I are given, an *instance* of the problem, Pr_I is given. Informally, a *parameterized decision problem* $Pr(k)$ (and its instances) can be defined by designating a certain input parameter k so that its complexity is an explicit function of this parameter. With this background, the following definition

is adapted from [18].

Definition 3.12 Let $Pr(k)_I$ be an instance of a parameterized decision problem $Pr(k)$ and Pr_I the instance of the non-parameterized version of $Pr(k)$. Then $Pr(k)$ is *(strongly uniformly) fixed-parameter tractable* if there is an algorithm that decides whether $Pr(k)_I$ is a yes-instance of $Pr(k)$ in time $f(k_s)\mathcal{O}(n^c)$ where n is the size of Pr_I , k_s is an integer parameter, c is a constant and f a recursive function. \square

In order to show that abduction over the well-founded semantics is fixed-parameter tractable, consider the decision problem of whether Q is contained in the abductive solution (Definition 2.8) to a finite, ground abductive framework, $\langle P, \mathcal{A}, I \rangle$. Then, since $C_{context}$ of Theorem 3.4 is not greater than $C_{abducibles}$, k_s can be set to $C_{abducibles}$ and $f(k_s)$ the summation, M , in Theorem 3.4 with $C_{context}$ replaced by $C_{abducibles}$. Furthermore, $size(dual((P \cup I), \mathcal{A}))$ is linear in the size of $P \cup A \cup I$ by Lemma A.5, so that when the maximal size of an abductive context is factored out, evaluation of Q requires a number of ABDUAL operations linear in the size of $P \cup A \cup I$.

The above considerations lead to the following Theorem.

Theorem 3.5 Let Q be a query to a finite ground abductive framework $\langle P, \mathcal{A}, I \rangle$, and let $C_{abducibles}$ be the size of the set \mathcal{A} . Then the problem of deciding whether Q is contained in an abductive solution to $\langle P, \mathcal{A}, I \rangle$ is fixed-parameter tractable with respect to $C_{abducibles}$.

There are, of course other means for parameterizing abduction over the well-founded semantics. For instance, an estimate of the maximum cardinality of contexts of abductive subgoals could be made via a suitably defined dependency graph, so that the input parameter $C_{abducibles}$ could be replaced by this parameter. Alternately, ABDUAL might be adapted so that a restriction were placed on the size of all abductive contexts. Such an approach might be relevant to using ABDUAL to solve model-based diagnosis problem, where attention was restricted to identifying single or double faults in the model.

4 Construction of Generalized Stable Models through ABDUAL

The three-valued abductive frameworks of Section 2 are not the only semantics used for abduction: Generalized Stable Models [22] provide an important alternative. In [7] it was shown that the abductive framework of Section 2 has the same expressive power as generalized stable models. In this section, we reformulate these results to show that ABDUAL can be used to evaluate abductive queries over generalized stable models. By allowing all positive literals to be inferred through abduction, ABDUAL can be used to construct partial stable interpretations (Definition 2.4). By choosing appropriate integrity constraints, these interpretations can be constrained to be consistent and total. We begin by adapting the concept of a generalized stable model to the terminology of Section 2.

Definition 4.1 [Generalized Partial Stable Interpretation and Model] Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework, with a scenario $\sigma = \langle P, \mathcal{A}, \mathcal{B}, I \rangle$. Then $M(\sigma)$ is a generalized partial stable interpretation of $\langle P, \mathcal{A}, I \rangle$ if

- $M(\sigma)$ is a partial stable interpretation of $\langle P \cup P_{\mathcal{B}} \cup I \rangle$; and

- \perp is false in $M(\sigma)$.

If in addition $M(\sigma)$ is an answer set of $\langle P \cup P_{\mathcal{B}} \cup I \rangle$, σ is a generalized stable model of $\langle P, \mathcal{A}, I \rangle$.
 \square

Generalized stable models can be computed by adding additional program rules, abducibles, and integrity constraints to abductive frameworks and computing the solution to these frameworks as per Definition 2.8.

Definition 4.2 Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework. Then let \mathcal{S} be the smallest set containing a new objective literal, abd_O , not in $literals(P \cup I \cup \mathcal{A})$ for each objective literal O in $literals(P \cup I \cup \mathcal{A})$. A literal formed over an element of \mathcal{S} is called a *shadow literal*. Let

$$R = O :- Body$$

be a rule in $P \cup I$. Then a *shadow rule* for R is a rule

$$R_{shadow} = O :- Body_{abd}$$

in which each literal of the form $not(O')$ in $Body$ is replaced by $not(abd_O')$. The shadow rules are denoted $Shadow(P)$ for a program P . Corresponding to these shadow rules are *shadow constraints* (I_{shadow}) of the form

$$\begin{aligned} \perp &:- O, not(abd_O). \\ \perp &:- not(O), abd_O. \end{aligned}$$

for each abd_O such that $not(abd_O) \in literals(Shadow(P))$.

The *consistency constraints* ($I_{consist}$) for $\langle P, \mathcal{A}, I \rangle$ consist of the shadow constraints along with integrity rules of the form

$$\perp :- O, not(O).$$

for $O \in literals(P \cup I)$.

The *totality rules* (I_{total}) for $\langle P, \mathcal{A}, I \rangle$ have the form

$$\begin{aligned} \perp &:- not(defined_O) \\ defined_O &:- O \\ defined_O &:- not(O) \end{aligned}$$

for each $O \in literals(P \cup I)$. \square

Example 4.1 Consider the abductive framework consisting of the program P_3 :

$$\begin{aligned} p &:- not\ q. \\ q &:- not\ p. \end{aligned}$$

with an empty set of abducibles and integrity constraints. In order to compute the partial stable interpretations of P_3 via abductive solutions, shadow rules must be added along with integrity and consistency constraints. For simplicity, we ignore coherency rules below. The shadow rules of P , $Shadow(P)$ are

$p \text{ :- not } abd_q.$
 $q \text{ :- not } abd_p.$

While the shadow constraints, I_{shadow} , include the rules

$\perp \text{ :- } p, \text{ not } abd_p$
 $\perp \text{ :- } q, \text{ not } abd_q$
 $\perp \text{ :- not } p, abd_p$
 $\perp \text{ :- not } q, abd_q$

and the consistency constraints, $I_{consist}$, include all instantiations of the schemata

$\perp \text{ :- } 0, \text{ not } 0$
 $\perp \text{ :- } 0, -0$

for $O \in \text{literals}(P \cup I)$. Let \mathcal{A} be the set $\{abd_O \text{ or } -abd_O | not(abd_O) \in \text{literals}(Shadow(P))\}$. Then the abductive framework

$$\langle (P \cup Shadow(P)), A, (I_{shadow} \cup I_{consist}) \rangle:$$

has solutions

$$\begin{aligned}\sigma_1 &= \{abd_q, abd_p\} \\ \sigma_2 &= \{-abd_q\} \\ \sigma_3 &= \{-abd_p\}\end{aligned}$$

These solutions correspond to the following generalized partial stable interpretations of $\langle P, \emptyset, \emptyset \rangle$, whose restrictions to the atoms of P are:

$$\begin{aligned}M(\sigma_1)|_{\{p,q\}} &= \emptyset \\ M(\sigma_2)|_{\{p,q\}} &= \{p\} \\ M(\sigma_3)|_{\{p,q\}} &= \{q\}\end{aligned}$$

Note that, in accordance with the definitions of Section 2, positive and negative objective literals are abduced, and coherency propagates negation from abduced objective literals to negative literals. In order to derive the generalized stable models of $\langle P, \emptyset, \emptyset \rangle$, the totality constraints of Definition 4.2 must also be added. In the above example, the totality constraints would prevent the first scenario, σ_1 , from being an abductive solution.

Example 4.1 illustrates the following theorem.

Theorem 4.1 *Let $F = \langle P, \mathcal{A}, I \rangle$ be an abductive framework, and $\sigma = \langle P, A, \mathcal{B}, I \rangle$ be an abductive scenario for F . Let $Shadow(P \cup I)$ be the set of shadow rules for $P \cup I$, and let I_{shadow} , $I_{consist}$, and I_{total} be the shadow, consistency, and totality constraints for F as in Definition 4.2. Let $\mathcal{A}_{shadow} = \{abd_O | not(abd_O) \in Shadow(P \cup I)\}$. Then*

1. $M(\sigma)$ is a generalized partial stable interpretation of $\langle P, A, I \rangle$ iff there exists an abductive solution

$$\sigma' = \langle (P \cup Shadow(P \cup I)), (A \cup \mathcal{A}_{shadow}), \mathcal{B}, (I \cup I_{shadow}) \rangle$$

such that $M(\sigma) = M(\sigma')$.

2. $M(\sigma)$ is a generalized stable model of $\langle P, A, I \rangle$ iff there exists an abductive solution

$$\sigma' = \langle (P \cup \text{Shadow}(P \cup I)), (A \cup A_{\text{shadow}}), \mathcal{B}, (I \cup I_{\text{shadow}} \cup I_{\text{consist}} \cup I_{\text{total}}) \rangle$$

such that $M(\sigma) = M(\sigma')$.

Proof: This result is straightforward from Definition 4.1 and the results of [7]. ◇

Theorem 4.1 has several implications. First, since the paraconsistent well-founded model of a program is a partial stable interpretation, use of the shadow program and constraints includes computation of the paraconsistent well-founded model as a special case. In addition, because Theorem 3.2 states that ABDUAL can be used for query evaluation to abductive frameworks based on *WFS*, ABDUAL can be used to compute queries to generalized partial stable interpretations and generalized stable models. The cost of this computation, of course, includes the cost of potentially evaluating shadow rules and the various additional integrity constraints. It is known that the problem of deciding the answer to a ground query to an abductive framework is NP-complete when the entailment method is based on the well-founded semantics [14], as is the problem of deciding whether an abductive framework has a generalized stable model. The lack of polynomial data complexity of ABDUAL for arbitrary abductive frameworks is therefore understandable, given the power of these frameworks. Finally, we note that computation of consistent answer sets can also be obtained via the transformation in Definition 4.2.

5 Discussion

5.1 A Meta-interpreter for ABDUAL and its Applications

Currently the ABDUAL system is implemented on top of the XSB System [34]. It consists of a preprocessor for generating the dual program, plus a meta-interpreter for the tabled evaluation of abductive goals, and is available from <http://www.cs.sunysb.edu/~tswift>. This meta-interpreter has the termination property of Theorem 3.3, but does not have the complexity property of Theorem 3.4. Work is currently being done in order to migrate into the XSB engine some of the tabling mechanisms of ABDUAL now taken care by the meta-interpreter, such as the CO-UNFOUNDED SET REMOVAL operation.

Psychiatric Diagnosis ABDUAL was originally motivated by a desire to implement psychiatric diagnosis [16]. Knowledge about psychiatric disorders is codified by *DSM-IV* [10] sponsored by the American Psychiatric Association. Knowledge in DSM-IV can be represented as a directed graph with positive links to represent relations from diagnoses to sub-diagnoses or to symptoms. These graphs also have negative links, called *exclusion* links that represent symptoms or diagnoses that must shown false in order to derive the diagnosis. The DSM-IV graph requires both abduction and non-stratified negation, as can be seen by considering the diagnosis of Adjustment Disorder ([10], pg. 626). One criterion for this diagnosis is

Once the stressor (or its consequences) has terminated, the symptoms do not persist for more than an additional 6 months.

Thus, to diagnose a patient as presently undergoing adjustment disorder, a physician must hypothesize about events in the future — a step naturally modeled with abduction. Adjustment disorder requires an exclusion criterion

The stress-related disturbance does not meet the criteria for another specific Axis I disorder and is not merely an exacerbation of a preexisting Axis I or Axis II disorder.

that admits the possibility of a loop through negation between adjustment disorder and another diagnosis. This can in fact occur, for instance with Alzheimer's Dementia ([10], pg. 142-143). If, as far as a physician can tell, a patient fulfills all criteria for adjustment disorder besides the above criterion, as well as all criteria for Alzheimer's (besides the criterion that the disturbance is not better accounted for by another disorder), the physician will essentially be faced with the situation:

The patient has an Adjustment Disorder if he does not have Alzheimer's Dementia, and has Alzheimer's Dementia if the patient does not have an Adjustment Disorder.

Use of abduction over DSM-IV must therefore handle non-stratified programs. The current user interface of the Diagnostica system (<http://medicinerules.com>) uses abduction in a simple but clinically relevant way to allow for hypothetical diagnosis: when there is not enough information about a patient for a conclusive diagnosis, the system allows for hypothesizing possible diagnosis on the basis of the limited information available.

Model-based Diagnosis ABDUAL has also been employed to detect specification inconsistencies in model-based diagnosis system for power grid failure [4]. Here abduction is used to attempt to abduce hypothetical physically possible events that might cause the diagnosis system to come up with a wrong diagnosis, violating the specification constraints. It is akin to model verification: one strains to abduce a model, comprised of abduced physical events, which attempts to make the diagnostic program inconsistent. If this cannot be done, the power grid can be certified to be correct. The attempt is conducted by trying to abduce hypothetical real world events which would lead to a proof of *falsum*, the atom reserved for the purpose of figuring in the heads of integrity constraints having the form of denials.

In this case, the application concerns a real electrical power grid network in Portugal, which is being monitored in real time by a pre-existing model-based logic programming diagnosis system (SPARSE)⁸ that receives time-stamped event report messages about the functioning or malfunctioning of the grid. The aim of our abductive application was to certify that a given expert system diagnosis module was provably correct with respect to foreseen physical events. To wit, the diagnosis logic program was executed under ABDUAL in order to establish that no sequence of (abduced) physically coherent events (i.e. monitoring messages) could be conducive to a diagnosis violating the (temporal) constraints expected of a sound diagnosis.

This approach proved to be feasible, though it required us to introduce a constructive negation implementation of ABDUAL, not yet reported elsewhere, because the abduced message events had to be time-stamped with temporally constrained conditions with variable parameters, and often these

⁸cf. <http://www.cim.isep.ipp.pt/Projecto-SATOREN/>

occurred under default negated literals (and hence the need for applying constructive negation on those variables), to the effect that no supervening event took place in some time related interval. The system, the application, and its use are described in detail in [4].

Four steps were involved in this process:

- Translation of the SPARSE rules into a syntactical form suitable for abduction.
- Preprocessing of the translated rules for use by our ABDUAL implementation.
- Obtaining abductive event solutions for diagnosis goals.
- Checking for physical consistency of the abductive solutions.

The most difficult and critical step was the first one, as the pre-existing SPARSE expert system rules had been written beforehand by their developers, with no abductive use in mind at all. Specific tools were developed to automate this step. The pre-existing ABDUAL implementation (comprising constructive negation) mentioned at step three (which required minimal adaptation), and the dualization preprocessor, mentioned at step two, both functioned to perfection. Step four was enacted by constructing tools to automate the analysis of the physical meaningfulness of the abduced solutions.

A number of open problems worthy of exploration remain in this class of problems, susceptible of furthering the use of the general abductive techniques employed.

Reasoning about Actions ABDUAL has been applied as well to model and reason about actions [3]. For this the ABDUAL system was integrated with Dynamic Logic Programming (DLP) Updates system [1].

DLP considers sequences of logic programs $P_1 \oplus P_2 \oplus \dots P_n$, whose intended meaning is the result of updating the logic program P_1 with the rules in P_2 , then updating the resulting knowledge base with \dots , and then updating the resulting knowledge base with P_n . In [1] a declarative semantics for DLP is presented. In order to ease the implementation of DLP, [1] also presents an alternative, equivalent, semantics which relies on a transformation of such sequences of programs into a single logic program in a meta-language. This transformation readily provides an implementation of DLP (obtainable via <http://centria.di.fct.unl.pt/~jja/updates>). For this implementation the use of tabling is of importance. In fact, the transformation relies on the existence of inertia rules for literals in the language, stating that some literal is true at some state if it was true before and is not overridden at that present state. Tabling is important for the efficiency of the implementation by avoiding repetition of computation for past states.

DLP has been used in applications for reasoning about actions [3]. In this setting actions are coded as logic programs updates which may have pre-conditions and post-conditions. For these applications the possibility of having programs with loops over negation is crucial. In fact, rules involved in such loops are used to model for instance unknown initial conditions and unknown outcomes of actions. For a concrete example, if one wants to state that initially it is not known whether or not the individual a was alive, one may write, in the first program P_1 , the rules:

$$\begin{aligned} \text{alive}(a) &\leftarrow \text{not}\neg\text{alive}(a) \\ \neg\text{alive}(a) &\leftarrow \text{not}\text{alive}(a) \end{aligned}$$

Reasoning about actions in a scenario is performed by a well-founded evaluation of the sequence of updated programs. Abductive reasoning is used for planning in the actions scenario. In fact, in this update setting, abducing updates (which code actions) in order to fulfill some goal of some future state amounts to plan which actions need to be executed in order to make that goal true. For this, a system with tabling, ability to deal with programs with loops over negation, and abduction was needed. ABDUAL includes all these ingredients, and was successfully employed for this purpose.

5.2 Comparisons with Other Methods

The use of dual programs to compute the well-founded semantics of normal programs was introduced in [27], but this method has several limitations compared with ABDUAL: it does not handle abduction or explicit negation; and it can have exponential complexity for some queries. Many of the definitions of ABDUAL are derived from SLG [5] (as reformulated in [33]) which computes queries to normal programs according to the well-founded semantics. For normal programs, ABDUAL shares the same finite termination and polynomial complexity properties as SLG. ABDUAL adds the capability to handle abduction (by adding abductive contexts to goals, modifying operations on forests to deal with such contexts, and by adding the ABDUCTION operation), adds the use of the dual transformation for extended programs and the CO-UNFOUNDED SET REMOVAL operation, but ABDUAL does not allow evaluation of a non-ground program as does SLG. Unfortunately, performance trade-offs of ABDUAL and SLG are not yet available, due to the lack of an engine-level implementation of the CO-UNFOUNDED SET REMOVAL operation of ABDUAL.

The main contribution of ABDUAL is its incorporation of abduction. We are not aware of any other efforts that have added abduction to a tabling method. Indeed, it is the use of tabling that is responsible for the termination and complexity results of Sections 3.4 and 4. Furthermore, ABDUAL evaluations are confluent in the sense that Theorem 3.2 holds for *any* ordering of applicable ABDUAL operations. The complexity and termination for *WFS* distinguishes ABDUAL from approaches such as the IFF proof procedure [15] and SLDNFA [9]. At the same time, these approaches do allow variables in rules which ABDUAL does not. The methods of [6] and [19] compute abductive explanation based on some form of two-valued rule completion for non-abducible predicates (the former based on Clark's completion, and the latter based on the so-called transaction programs). This is similar to our use of the dual program⁹. In both methods, abductive explanations are computed by using the only-if part of the completion in a bottom-up fashion. However, both methods have a severe restriction on the class of programs: they apply generally only to acyclic programs. This restriction is due to their being based on completion so that from an ABDUAL perspective, these methods do not require operational analogs to the DELAYING and SIMPLIFICATION operations to evaluate unfounded sets of objective literals, or the CO-UNFOUNDED SET REMOVAL operation to evaluate co-unfounded sets. The pay-off of

⁹Note that the dual for non-abducible predicates in acyclic programs is the same as the completion.

adding these operations is that ABDUAL is based on the well-founded semantics, and does not impose any restriction on cycles in programs.

The restriction on cycles is also not imposed by methods based on the stable models semantics, such as [30, 31, 23]. As ABDUAL, the method of [30] also requires a prior program transformation. In this case, an abductive programs is translated into a normal logic program, such that the stable models of the latter correspond to the abductive solutions of the former. This method has some drawbacks. Most importantly, by doing so, one may obtain abductive solutions with atoms that are not relevant for the abductive query. To avoid this drawback, in [31] the method is improved by incorporating a top-down procedure to determine the relevance of the abducible to the query. The ACLP system of [23] is based on Generalized Stable Models, but it also integrates in a single framework abduction and constraint programming. Again the complexity results for *WFS*, when compared to that for the stable models semantics, distinguishes ABDUAL from these approaches.

We have shown in this article how ABDUAL can be mustered to compute Generalized Stable Models, and thus Stable Models in particular. Some words are in order on comparing it to other Stable Model implementations, such as DLV [13], and S-Models [32]. These implementations are specialized toward Stable Model evaluation, and are restricted to finite ground programs without functional symbols, though some preprocessors can help to do the grounding where possible and domain information is available. Naturally, their efficiency for the specific purpose of computing Stable Models is better than that of a general procedure like ABDUAL, even though the complexity remains the same.

Abduction can also be carried out by those specialized implementations by means of known program transformations, such as the ones shown in [24, 29]. Though one common problem to those approaches is that, because of the non-relevancy character of Stable Models, and also of abducibles being two-valued in them, all possible (non-minimal) abductions are potentially generated, and not just those relevant for a top goal.

With respect to stratified programs, where the well-founded and stable models semantics coincide, ABDUAL is able to deal with function symbols and non-ground programs in infinite domains, and perform demand driven abduction. Moreover, if abduction is not after all required, then the complexity, we have seen, remains polynomial, and no unnecessary abductions are made, in contradistinction to the two-valued approach, which requires for all abducibles to be abducted either as true or as false.

In summary, the two approaches are designed for different purposes, and each should excel in its own territory. Protracted attempts to have ABDUAL compute a relevant residual program that would be passed on to an implementation of stable models have failed, as most of the work ends up having to be done on the ABDUAL side, without the desired sharing of specialized effort.

Generalizing ABDUAL to Programs with Variables Generalizing ABDUAL for non-ground covered programs¹⁰ with ground queries is not a difficult task: as in Clark's completion, consider rule heads with free variables, and explicitly represent unifications in the body; the dual is then obtained from these rules as usual, where the negation of $=$ is \neq . Allowing non-ground queries in covered

¹⁰A program is covered iff all variables appearing in the body of rules also appear in the corresponding head.

programs can be obtained by considering as abducibles all terms of the form $X \setminus = T$, and by adding an appropriate method for verifying consistency of sets of such inequalities. Such a method could greatly benefit from an integration of ABDUAL with constraint programming, where the consistency of the inequalities would be checked by a constraint solver. The integration of abduction and constraint programming, as is already done for other systems (viz. SLDNFAC [26] and ACLP [23]), is in our research agenda.

The most difficult step in order to fully generalize ABDUAL to deal with non-ground programs is to abandon the restriction of covered programs. This is so because free variables in the body of program rules introduce universally quantified variables in the body of rules in the dual program — a problem similar to that of floundering in normal programs. Work is underway to generalize ABDUAL to deal with non-ground non-covered programs using constructive negation methods.

A practical advantage of ABDUAL is that it allows the easy propagation of abducibles through both positive and negative literals. As an abductive answer is returned to an abductive subgoal, contexts can be immediately checked for consistency, regardless of whether the subgoal is positive or negative, and regardless of how many levels of negation were needed to produce the answer.

Acknowledgements This work was partially supported by NSF grants CCR-9702581, EIA-97-5998, and INT-96-00598. The authors also thank PRAXIS XXI projects FLUX and FLAD-NSF project REAP for their support. Finally, the presentation of this paper was considerably improved with the help of detailed comments from an anonymous reviewer.

References

- [1] J. Alferes, J. Leite, L. M. Pereira, H. Przymusinska, and T. Przymusinski. Dynamic updates of non-monotonic knowledge bases. *Journal of Logic Programming*, 45(1-3):43–70, 2000.
- [2] J. Alferes, L. M. Pereira, and T. Swift. Well-founded abduction via tabled dual programs. In *International Conference on Logic Programming*, pages 426–440, 1999.
- [3] J. J. Alferes, J. A. Leite, L. M. Pereira, and P. Quaresma. Planning as abductive updating. In D. Kitchin, editor, *Procs. of AISB 2000*, 2000.
- [4] J.F. Castro and L.M. Pereira. Abductive validation of a power-grid diagnoser, 2002. English version of Master’s thesis Verificação Abdutiva de um Sistema de Diagnóstico Baseado em Regras. Available at <http://www.cs.sunysb.edu/~tswift>.
- [5] W. Chen and D. S. Warren. Tabled Evaluation with Delaying for General Logic Programs. *Journal of the ACM*, 43(1):20–74, January 1996.
- [6] L. Console, D. Dupré, and P. Torasso. On the relationship between abduction and deduction. *Journal of Logic and Computation*, 1(5):661–690, 1991.

- [7] C. Damásio and L. M. Pereira. Abduction over 3-valued extended logic programs. In *International Conference on Logic Programming and Non-Monotonic Reasoning*, pages 29–42. Springer-Verlag, 1995. LNAI 1265.
- [8] C. V. Damásio and L. M. Pereira. A survey on paraconsistent semantics for extended logic programas. In D.M. Gabbay and Ph. Smets, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, volume 2, pages 241–320. Kluwer Academic Publishers, 1998.
- [9] M. Denecker and D. De Schreye. SLDNFA: An abductive procedure for normal abductive programs. *J. Logic Programming*, 34(2):111–167, 1998.
- [10] *Diagnostic and Statistical Manual of Mental Disorders*. American Psychiatric Association, Washington, DC, 4th edition, 1994. Prepared by the Task Force on DSM-IV and other committees and work groups of the American Psychiatric Association.
- [11] J. Dix and T. Swift. Linear subclasses of the well-founded semantics, 2002. In preparation.
- [12] R. G. Downey and M.R. Fellows. Fixed parameter tractability and completeness I: Basic results. *SIAM J. Computing*, 24:873–921, 1995.
- [13] T. Eiter, N. Leone, C. Mateis, G. Pfeifer, and F. Scarcello. A deductive system for nonmonotonic reasoning. In J. Dix, U. Furbach, and A. Nerode, editors, *Logic Programming and Nonmonotonic Reasoning*, volume 1265 of *LNAI*. Springer, 1997.
- [14] Thomas Eiter, Georg Gottlob, and Nicola Leone. Abduction From Logic Programs: Semantics and Complexity. *Theoretical Computer Science*, 189(1-2):129–177, December 1997.
- [15] T. Fung and R. Kowalski. The IFF proof procedure for abductive logic programming. *Journal of Logic Programming*, 33(2):151–165, 1997.
- [16] J. Gartner, T. Swift, A. Tien, L. M. Pereira, and C. Damásio. Psychiatric diagnosis from the viewpoint of computational logic. In Lloyd et. al, editor, *International Conference on Computational Logic*, pages 1362–1376. Springer-Verlag, 2000. LNAI 1861.
- [17] M. Gelfond and V. Lifshitz. Logic programs with classical negation. In *International Conference on Logic Programming*, pages 579–597. MIT Press, 1990.
- [18] G. Gottlob, F. Scarcello, and M. Sideri. Fixed parameter complexity in AI and nonmonotonic reasoning. In *International Conference on Logic Programming and Non-Monotonic Reasoning*, pages 1–18, 1999.
- [19] K. Inoue and C. Sakama. Computing extended abduction through transaction programs. *Annals of Mathematics and Artificial Intelligence*, 25(3-4):339–367, 1999.
- [20] J.F.Castro. Verificação abdutiva de um sistema de diagnóstico baseado em regras (Abductive verification of a rule based diagnosis system). Master’s thesis, Universidade Nova de Lisboa, December 2000. In Portuguese. Available at <http://www.cs.sunysb.edu/~tsswift>.

- [21] A. Kakas, R. Kowalski, and F. Toni. Abductive logic programming. *Journal of Logic and Computation*, 2(6):719–770, 1993.
- [22] A. Kakas and P. Mancarella. Generalized stable models: A semantics for abduction. In *ECAI*, pages 385–391. Morgan-Kaufmann, 1990.
- [23] A. C. Kakas and C. Mourlas. ACLP: Flexible solutions to complex problems. In *International Conference on Logic Programming and Non-Monotonic Reasoning*, pages 387–399. Springer, 1997. LNAI 1265.
- [24] V. Lifshitz and H. Turner. From disjunctive programs to abduction. In J. Dix, L. M. Pereira, and T. Przymusiński, editors, *Nonmonotonic Extensions of Logic Programming*, volume 927 of *LNAI*, pages 23–42. Springer, 1995.
- [25] J. W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, Berlin Germany, 1984.
- [26] B. Van Nuffelen and M. Denecker. Problem solving in ID-logic with aggregates: some experiments. In *8th Int. Workshop on Non-Monotonic Reasoning (NMR'2000)*, pages 9–11, 2000. Available via www.cs.kuleuven.ac.be.
- [27] L. M. Pereira, J. Aparício, and J. Alferes. Derivation procedures for extended stable models. In *International Joint Conference on Artificial Intelligence*, 1991.
- [28] T.C. Przymusiński. Every logic program has a natural stratification and an iterated least fixed point model. In *ACM Principles of Database Systems*, pages 11–21. ACM Press, 1989.
- [29] C. Sakama and K. Inoue. Abductive logic programming and disjunctive logic programming: Their relationship and transferability. *Journal of Logic Programming*, 44(1-3):71–96, 2000.
- [30] K. Satoh and N. Iwayama. Computing abduction by using the TMS. In *International Conference on Logic Programming*, pages 505–518. MIT Press, 1991.
- [31] K. Satoh and N. Iwayama. Computing abduction by using TMS and top-down expectation. *J. Logic Programming*, 44(1-3):179–206, 2000.
- [32] *Smodels*, 2003. Home page: <http://www.tcs.hut.fi/Software/smodels>.
- [33] T. Swift. A new formulation of tabled resolution with delay. In *Recent Advances in Artificial Intelligence*, pages 163–177. Springer-Verlag, 1999. Available at <http://www.cs.sunysb.edu/~tsswift>.
- [34] The XSB Logic Programming System version 2.5, 2002. <http://xsb.sourceforge.net>.
- [35] A. van Gelder, K.A. Ross, and J.S. Schlipf. Unfounded sets and well-founded semantics for general logic programs. *Journal of the ACM*, 38(3):620–650, 1991.

A Appendix: Proofs of Theorems

A.1 Proof of Theorems in Section 3.3

Theorem 3.1 *For a ground program P and empty set of abducibles, let $\text{dual}(P, \emptyset)$ be the dual program formed by applying the dual transformation to P . Then for any literal $L \in \text{literals}(P)$,*

$$L \in \text{WFS}(P) \iff L \in \text{lfp}(\omega_d^{\text{dual}(P, \emptyset)})$$

Proof:

As mentioned in Section 3, if P is a countable set of rules, each with a finite number of literals in their bodies, $\text{dual}(P, \emptyset)$ will be also.

The inner fixed point of $\omega_d^{\text{dual}(P, \emptyset)}$ depends on two operators: $Td_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$ and $Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$ (Definition 3.10). $Td_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$ is essentially the same as the inner fixed point operator, $Tx_{\mathcal{I}}^P$ (Definition 2.1) used to construct $\text{WFS}(P)$. Combining this observation with the initial sets used to construct the fixed points of Definitions 2.2 and 3.11, the proof can be reduced to showing that for an element O in $\text{objective_literals}(P)$,

$$O \in \text{gfp}(Fx_{\mathcal{I}}^P(\text{objective_literals}(P))) \iff \text{not}(O) \in \text{gfp}(Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{S}_{\text{dual}(P, \emptyset)}))$$

If this is so, since $Tx_{\mathcal{I}}^P$ equals $Td_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$ restricted to the objective literals in P , then it is a trivial induction on the operators of Definitions 2.2 and 3.11 to show that ω_{ext}^P equals $\omega_d^{\text{dual}(P, \emptyset)}$ restricted to the literals in P .

Note that the initial set used to construct the fixed point of $Fx_{\mathcal{I}}^P$ in Definition 2.2 is $\text{objective_literals}(P)$, while in Definition 3.11 the initial set used for $Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$ is

$$\mathcal{S}_{\text{dual}(P, \mathcal{A})} = \{\text{not}(O) \mid \text{not}(O) \in \text{literals}(\text{dual}(P, \mathcal{A}))\}$$

and thus by the dual transformation an objective literal O occurring in P , O is in $\text{objective_literals}(P)$ iff $\text{not}(O)$ is in $\mathcal{S}_{\text{dual}(P, \mathcal{A})}$.

(\Rightarrow) We first prove that for $O \in \text{objective_literals}(P)$:

$$O \in \text{gfp}(Fx_{\mathcal{I}}^P(\text{objective_literals}(P))) \Rightarrow \text{not}(O) \in \text{gfp}(Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{S}_{\text{dual}(P, \emptyset)}))$$

Induction is on the number n of applications of $Fx_{\mathcal{I}}^P$ in constructing the fixed point. The base case, where $n = 0$ was handled in the previous paragraph, so consider first the case in which n is greater than 0. In other words, the two operators have been applied $n - 1$ times, with applications of $Fx_{\mathcal{I}}^P$ producing the set \mathcal{O}_1 and applications of $Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$ producing the set \mathcal{O}_2 . It remains to prove that for the n^{th} application:

$$O \in Fx_{\mathcal{I}}^P(\mathcal{O}_1) \Rightarrow \text{not}(O) \in Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{O}_2)$$

Suppose an objective literal O is in $Fx_{\mathcal{I}}^P(\mathcal{O}_1)$. Then either (1) $\text{conj}_E(O) \in \mathcal{I}$; or for every rule r_j^O of the form $O :- L_1, \dots, L_n$ in P , there is a literal $L_{j,i}$, $1 \leq i \leq n$, such that either (2a) $\text{conj}_D(L_{j,i}) \in \mathcal{I}$ or (2b) $L_{j,i}$ is in \mathcal{O}_1 . Consider each of these cases in turn.

1. For the first case, if $\text{conj}_E(O) \in \mathcal{I}$ then, by the rule $\text{not}(O) \text{ :- } \text{conj}_E(O)$ in $\text{dual}(P, \emptyset)$, $\text{not}(O)$ belongs to $Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{O}_2)$.
2. For the second case, consider a witness of unusability, $L_{j,i}$ for a rule r_j^O for O . (2a) Suppose first that $\text{conj}_D(L_{j,i}) \in \mathcal{I}$. Then there is a folding rule $\text{not}(\text{fold}_j^b O) \text{ :- } \text{conj}_D(L_{j,i})$ constructed by the dual transformation of Definition 3.1, so that $\text{not}(\text{fold}_j^b O)$ will be included in $Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{O}_2)$. (2b) Alternately, if $L_{j,i} \in \mathcal{O}_1$, then $\text{not}(L_{j,i}) \in \mathcal{O}_2$ so that $\text{not}(L_{j,i})$ and all literals in the heads of the folding rules that depend on it will be regenerated. In either case (2a) or (2b), by monotonicity of the operators, if $O \in Fx_{\mathcal{I}}^P(\mathcal{O}_1)$, then $O \in \mathcal{O}_1$, and by the induction hypothesis, $\text{not}(O) \in \mathcal{O}_2$. By definition of $Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$, for this to happen each rule for O in P must have a witness of unusability. This means that the literal $\text{not}(\text{fold}_k^a O_k)$ must also be in \mathcal{O}_2 for each rule k for O in P . Thus each of the literals $\text{not}(\text{fold}_l^b O_l)$ and $\text{not}(\text{fold}_k^b O_k)$ will be regenerated, so that $\text{not}(O) \in Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{O}_2)$.

(\Leftarrow) To complete the proof we need to show that for $O \in \text{objective_literals}(P)$:

$$O \in \text{gfp}(Fx_{\mathcal{I}}^P(\text{objective_literals}(P))) \Leftarrow \text{not}(O) \in \text{gfp}(Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{S}_{\text{dual}(P, \emptyset)}))$$

The proof of this statement is complicated by the fact that, due to folding rules, there is no exact correspondance between the iteration in which an objective literal is removed from a set by $Fx_{\mathcal{I}}^P$ and when its analog is removed by $Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}$. Accordingly, we restate this case as

$$O \notin \text{gfp}(Fx_{\mathcal{I}}^P(\text{objective_literals}(P))) \Rightarrow \text{not}(O) \notin \text{gfp}(Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{S}_{\text{dual}(P, \emptyset)}))$$

Consider a literal O' in

$$\text{objective_literals}(P) - \text{gfp}(Fx_{\mathcal{I}}^P(\text{objective_literals}(P)))$$

Because $Fx_{\mathcal{I}}^P$ is monotonic and continuous, there is an application n such that $O' \in \text{gfp}(Fx_{\mathcal{I}}^{P^{n-1}}(\text{objective_literals}(P)))$ and $O' \notin \text{gfp}(Fx_{\mathcal{I}}^{P^n}(\text{objective_literals}(P)))$. We call n the *removal level* of O .

Thus we prove by induction on n that any objective literal with removal level n is not in $\text{gfp}(Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{S}_{\text{dual}(P, \emptyset)}))$. The base case (where $n = 0$) that for $O \in \text{objective_literals}(P)$, $O \in \text{objective_literals}(P) \iff \text{not}(O) \in \mathcal{S}_{\text{dual}(P, \emptyset)}$ was handled above. Assume that the property holds for objective literals in P whose removal level is less than n , and that the set produced by the first $n - 1$ applications of $Fx_{\mathcal{I}}^P$ is \mathcal{O}_1 .

Consider a literal O with removal level n . By Definition 2.1, O will not persist after the n th application of $Fx_{\mathcal{I}}^P$ iff $\text{conj}_E(O) \notin \mathcal{I}$ and there exists a rule r_O^k with body L_1, \dots, L_m such that for all $1 \leq i \leq m$, $\text{conj}_D(L_i) \notin \mathcal{I}$ and $L_i \notin \mathcal{O}_1$.

First note that $\text{conj}_E(O) \notin \mathcal{I}$ means that an axiom of coherence will not be used to derive $\text{not}(O)$ in $\text{gfp}(Fd_{\mathcal{I}}^{\text{dual}(P, \emptyset)}(\mathcal{S}_{\text{dual}(P, \emptyset)}))$. Next, by Definition 3.1, there is a folding rule

$$\text{not}(\text{fold}_k^b O) \text{ :- } \text{conj}_D(L_i)$$

for each L_i in the body of r_O^k . Note that $conj_D(L_i)$ is not in \mathcal{I} (otherwise O would not have a removal level). If $L_i \notin \mathcal{O}_1$, then the removal level of L_i is less than n so by the induction assumption, there will be some application of $Fd_{\mathcal{I}}^{dual(P, \emptyset)}$ that does not regenerate $conj_D(L_i)$. Accordingly, there will be an application of $Fd_{\mathcal{I}}^{dual(P, \emptyset)}$ in which $not(fold_k^b O)$ is not regenerated via its i th clause. For O to have a removal level, each literal in the body of r_O^k must have similar conditions to L_i , and $not(fold_k^b O)$ is not regenerated. By Definition 3.1, regeneration of $not(fold_k^b O)$ is required to regenerate $not(fold_k^a O)$. In future iterations, then, $not(fold_j^a O)$ will not be regenerated for $1 \leq j \leq k$, so that O will not be in $gfp(Fd_{\mathcal{I}}^{dual(P, \emptyset)}(\mathcal{S}_{dual(P, \emptyset)}))$.

◇

In addition to Theorem 3.1, several lemmas and definitions will be needed to prove the correctness of ABDUAL. One of these is a simpler definition of the dual transformation, which is convenient to use in the proofs. This definition was implicitly used in Examples 3.1, 3.3 and 3.4.

Definition A.1 [Unfolded dual Program] Let P be a ground extended program, and \mathcal{A} a (possibly empty) set of finite abducibles. Then the *unfolded dual transformation* creates a dual program $udual(P, \mathcal{A})$, defined as union of P with the smallest program containing the sets of rules R_1 and R_2 defined as follows:

1. If P contains a rule with non-empty body

$$\begin{aligned} O &:- L_{1,1}, \dots, L_{1,n_1} \\ &: \\ O &:- L_{m,1}, \dots, L_{m,n_m} \end{aligned}$$

Then, R_1 contains the rules

$$not(O) :- conj_D(L_{1,j_1}), \dots, conj_D(L_{m,j_m}).$$

for each j , $1 \leq j_i \leq i_m$, $1 \leq i \leq m$, and where $conj_D(L)$ represents the default conjugate of L .

2. Otherwise, if $not(O)$ is in $literals(P)$, but there is no rule with head O in P , then R_1 contains the rule $not(O) :- \mathbf{t}$.
3. R_2 consists of *axioms of coherence* that relate explicit and default negation, defined as:

$$not(O) :- conj_E(O)$$

For each objective literal $not(O)$ in either $literals(P \cup R_1)$ or \mathcal{A} .

□

Because the dual transformation of Definition 3.1 differs from that of Definition A.1 only insofar as no folding rules are defined, it is straightforward to see that they are equivalent with respect to $literals(P)$ for finite ground programs. While Definition A.1 is simpler than Definition 3.1 and useful

for correctness results, Definition 3.1 is necessary for proving correctness of programs with an infinite number of rules and for the complexity results that follow.

We next make explicit the relation between unfounded sets of objective literals in P and co-unfounded sets of answers in the dual program. In order to do so, we present the definition of an interpretation induced by the state of an ABDUAL evaluation, and of the delay dependency graph.

Definition A.2 Given an ABDUAL forest \mathcal{F} , the *interpretation induced by \mathcal{F}* , or $\mathcal{I}_{\text{induced}(\mathcal{F})}$ is defined as the smallest interpretation containing

- A literal O for each unconditional answer node $\langle O, \emptyset \rangle \text{ :- } |$ in \mathcal{F} ;
- A literal $\text{not}(O)$ for each objective literal $O \in \mathcal{F}$ such that the tree for O is completely evaluated in \mathcal{F} and contains no answers.

□

The definition of a delay dependency graph is convenient for several of the following proofs.

Definition A.3 [Delay Dependency Graph] Let \mathcal{F} be an ABDUAL forest. A goal S_1 has a *direct delay dependency* on a goal S_2 in \mathcal{F} iff S_2 is contained in the *DelayList* of an answer in the tree T for S_1 . The *delay dependency graph* of \mathcal{F} , $\text{DDG}(\mathcal{F})$, is a directed graph $\{V, E\}$ in which V is the set of root goals for trees in \mathcal{F} and $(S_i, S_j) \in E$ iff S_i has a direct delay dependency on S_j . □

Also for convenience, if \mathcal{S} is a co-unfounded set of answers,

$$\text{heads}(\mathcal{S}) = \{S \mid \langle S, \text{Context} \rangle \text{ :- } DL \mid DL \in \mathcal{S}\}$$

Part (1) of Lemma A.1 relates a co-unfounded set of answers (Definition 3.6) obtained in the ABDUAL evaluation of a ground extended program to an unfounded set of objective literals (Definition 3.4) in the well-founded semantics (with explicit negation). Intuitively, part (2) ensures that when an ABDUAL forest can be constructed to capture an interpretation \mathcal{I} of a ground extended program, then the forest can be extended to determine the truth value of each objective literal in the unfounded sets with respect to \mathcal{I} .

Lemma A.1 Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework in which \mathcal{A} and I are empty.

1. Let \mathcal{F} be a forest in a ABDUAL evaluation \mathcal{E} of a query Q to $\langle P, \mathcal{A}, I \rangle$, and S^A be a co-unfounded set of answers in \mathcal{F} . Then there is a minimal unfounded set of objective literals S^O for P in $\mathcal{I}_{\text{induced}(\mathcal{F})}$ such that $\{S_i \mid \text{not}(S_i) \in \text{heads}(S^A)\} \subseteq S^O$.
2. Let S^O be a minimal unfounded set of objective literals for P with respect to an interpretation \mathcal{I} and let \mathcal{F} be an ABDUAL forest of any query to $\langle P, \mathcal{A}, I \rangle$ such that
 - (a) for each $S_i \in S^O$, there is a tree for $\text{not}(S_i)$ in \mathcal{F} ;
 - (b) for all $L \in I$, if $\langle L, \emptyset \rangle$ is the root of a tree in \mathcal{F} , then $L \in \mathcal{I}_{\text{induced}(\mathcal{F})}$;
 - (c) no ABDUAL operations are applicable to \mathcal{F} .

Then, for each $S_i \in S^O$, there will be an unconditional answer for each $\text{not}(S_i)$ in \mathcal{F} .

Proof: For simplicity of presentation, we first restrict our attention to finite programs, in which class the dual program, $\text{dual}(P, \emptyset)$, formed via Definition A.1 can be used. We then indicate how the arguments can be extended to the dual form of Definition A.1 which is necessary for infinite programs.

1. Let $N_{\text{leaf}} = \langle \text{not}(S_i), \emptyset \rangle \text{ :- } DL$ be an answer in some co-unfounded set of answers S^A in \mathcal{F} . By clause 3 of Definition 3.6, DL must be non-empty. By the construction of Definition A.1 and by the definitions of ABDUAL it can easily be seen that the node $\langle \text{not}(S_i), \emptyset \rangle \text{ :- } DL$ is a descendant of a non-root node $N = \langle \text{not}(S_i), \emptyset \rangle \text{ :- } L_1, \dots, L_n$ produced by application of a PROGRAM CLAUSE RESOLUTION operation of a rule r_u in $\text{dual}(P, \emptyset)$ (this situation is presented schematically in Figure 8) r_u is constructed so that for each rule r_j with head S_i in P , L_j is a default conjugate of some literal in the body of r_j . Now we consider two classes of the literals L_1, \dots, L_n in the (*GoalList* of the) node N : those that are contained in the *DelayList* of N_{leaf} and those that are not, and consider the latter first.




Figure 8: Schematic portion of an ABDUAL tree

- (a) Literals that are in N but not in the *DelayList* of N_{leaf} . By assumption, since the set of abducibles \mathcal{A} is empty, there can be no ABDUCTION operations in \mathcal{E} ; thus a literal L_i in N but not in N_{leaf} can have been resolved away either through ANSWER CLAUSE RESOLUTION, or through DELAYING and subsequent SIMPLIFICATION. Note that L_i cannot be created through a direct application of CO-UNFOUNDED SET REMOVAL as that operation acts only on leaf nodes, and removes all elements in a *DelayList*, contradicting the fact that DL is empty. In other words, L_i could be resolved away only via an ANSWER CLAUSE RESOLUTION or SIMPLIFICATION operation, and these operations would be possible only if L_i were in the interpretation induced by a previous forest, say \mathcal{F}_i . Furthermore, since no ABDUAL operation in Definition 3.9 allows operations on an unconditional answer node, all answers in \mathcal{F}_i will persist in \mathcal{F} , so that L_i must be in $\mathcal{I}_{\text{induced}(\mathcal{F})}$. It is then straightforward from Definitions A.1 and 3.4 that if $L_i \in \mathcal{I}_{\text{induced}(\mathcal{F})}$, then its conjugate will form a witness of unusability for some rule for S_i in P .
- (b) Literals that are in N and in the *DelayList* of N_{leaf} . Because $N_{\text{leaf}} \in S^A$, each literal, $L_i \in DL$ is such that $\langle L_i, C_i \rangle \text{ :- } DL_i \in S^A$, by Definition 3.6.

Taken together, the above two cases show that, given an conditional answer $\langle \text{not}(S_i), \emptyset \rangle :- DL \in S^A$, a witness of unsuitability can be obtained for every rule r_i for S_i in P . This can occur because r_i contains some literal $\text{conj}_D(L)$ that is true in $\mathcal{I}_{\text{induced}(\mathcal{F})}$, corresponding to condition (1) of Definition 3.4; or it can occur because the literal belongs to $\text{heads}(S^A)$, corresponding to condition (2) of Definition 3.4. Thus, in S^A , a witness of unusability can be obtained for every rule for S_i , so that $\{S_i | \text{not}(S_i) \in \text{heads}(S^A)\}$ is an unfounded set for $\mathcal{I}_{\text{induced}(\mathcal{F})}$ in P .

2. Since S^O is a minimal unfounded set of objective literals, no objective literal O_1 in \mathcal{I} can have its default conjugate in S^O (otherwise, $S^O - \{O_1\}$ would also be minimal). Let S_i be an objective literal in S^O , and let $L_i, 1 \leq i \leq n$, be witnesses of unusability for each of its (n) rules. By the construction of Definition A.1, there is a rule $r_{S_i} = \text{not}(S_i) :- \text{conj}_D(L_1), \dots, \text{conj}_D(L_n)$ in $\text{dual}(P, \emptyset)$. By assumption (a), \mathcal{F} contains a tree for $\text{not}(S_i)$. By assumption (c) and Definition 3.9, this tree must have children, otherwise there would be PROGRAM CLAUSE RESOLUTION operations applicable for \mathcal{F} . In particular, there must be a node $N' = \langle \text{not}(S_i), \emptyset \rangle :- \text{conj}_D(L_1), \dots, \text{conj}_D(L_n)$, that is a child of the root node for $\text{not}(S_i)$.

Now consider each $\text{conj}_D(L_i)$. Since L_i is a witness of unusability for S^O , L_i is either be false in \mathcal{I} or unfounded but not false in \mathcal{I} .

- (a) If L_i is false in \mathcal{I} , then by assumption (b) $\text{conj}_D(L_i)$ is true in $\mathcal{I}_{\text{induced}(\mathcal{F})}$, and either an ANSWER CLAUSE RESOLUTION or SIMPLIFICATION (after previous DELAYING) will be applicable, by assumption (c) these will have been performed and there will be a descendant of N' in which $\text{conj}_D(L_i)$ is resolved away.
- (b) Next, consider the case in which L_i is unfounded but not false in \mathcal{I} , and by assumption is in S^O . Because L_i is in an unfounded set it is a positive literal, and $\text{conj}_D(L_i)$ can be written as $\text{not}(L_i)$. Consider the node N' mentioned above: $\langle \text{not}(S_i), \emptyset \rangle :- \text{conj}_D(L_1), \dots, \text{conj}_D(L_n)$ for which $\text{not}(L_i)$ is a body literal. For $1 \leq i \leq n$, each $\text{conj}_D(L_i)$ is either in $\mathcal{I}_{\text{induced}(\mathcal{F})}$ (by assumption (b)) and resolved away by (assumption (c)); or its default conjugate is unfounded in \mathcal{I} and thus a negative literal, so that there is a node $\text{not}(S_i) :- DL$ in a previous forest $\mathcal{F}_{\text{prev}}$ of \mathcal{F} in \mathcal{E} , and $\text{not}(L_i) \in DL$ (by assumption (b)).

We must show that a CO-UNFOUNDED SET REMOVAL operation was applicable that removed $\text{not}(L_i)$ from DL . For this to happen, we must show that $\text{not}(L_i)$ was part of a co-unfounded set of answers for some previous forest $\mathcal{F}_{\text{prev}}$ in \mathcal{E} . Condition (2) of the definition of a co-unfounded set of answers (Definition 3.6) is trivially satisfied since the set of abducibles is empty, while condition (1) of Definition 3.6, stating that goals in the co-unfounded set of answers be completely evaluated, is satisfied by assumption (c). Condition (3) of Definition 3.6 remains, and we must show that there is a co-unfounded set of answers in $\mathcal{F}_{\text{prev}}$ containing an answer, $N_{L_i} = \text{not}(L_i) :- DL_i$ for $\text{not}(L_i)$. We begin by showing that N_{L_i} exists. Assumption (c) ensures that a tree for $\text{not}(L_i)$ exists in \mathcal{F} (by the argument above, $\text{not}(L_i)$ was selected, and assumption (c) ensures that a NEW SUBGOAL is performed when $\text{not}(L_i)$ was selected). Furthermore, assumption (c) implies that there are no applicable

ABDUAL operations for this tree. Furthermore, DL_1 contains all literals of a rule in the unfolded dual of P that gave rise to N_L , but which are not themselves in \mathcal{I} . Extending this argument for all elements in the transitive closure of $\text{not}(L_i)$ in the delay dependency graph of N (Definition A.3), shows that L_i is contained in an unfounded set of objective literals. Thus a CO-UNFOUNDED SET REMOVAL operation was applicable to \mathcal{F}_{prev} which made the answer for some $\text{not}(S_j)$, $S_j \in S'$ unconditional, and it can be easily seen that this operation made further SIMPLIFICATION operations applicable based on the unconditional answer for $\text{not}(S_j)$. Furthermore, by Definition 3.9 each SIMPLIFICATION operation for literal L and conditional answer $Ans \in S$ made applicable after the CO-UNFOUNDED SET REMOVAL operation will remain applicable until L is removed from the *DelayList* of Ans . Because the set of unconditional answers for a forest only grows monotonically, the statement holds.

Extending the proofs to the dual of infinite programs means that the dual transformation of Definition 3.1 must be used. In the case of infinite programs, there may be an infinite number of witnesses of unusability for a given objective literal O and interpretation \mathcal{I} . Unlike the situation presented above for finite programs, the witnesses of unusability may be distributed among different folding trees. With this complication the argument for part 1 of the lemma can be straightforwardly extended to the situation where folding rules are used. In part 2, note that assumptions 2a and 2b together imply that if O is an objective literal in P , then \mathcal{F} will also contain the appropriate folding literals for $\text{not}(O)$. Again, the argument can be straightforwardly extended to this new situation. \diamond

Lemma A.2 *Let $\langle P, \emptyset, \emptyset \rangle$ be an abductive framework. Let \mathcal{E} be a ABDUAL evaluation of a query Q against $\text{dual}(P, \emptyset)$, whose final forest is \mathcal{F}_β . Finally, let $WFS(P)|_{\mathcal{E}_\beta}$ denote the well-founded model of P restricted to goals in \mathcal{F}_β . Then*

$$\mathcal{I}_{induced(\mathcal{F}_\beta)} = WFS(P)|_{\mathcal{F}_\beta}$$

Proof:

(Sketch) Given Theorem 3.1 this is equivalent to showing that

$$\mathcal{I}_{induced(\mathcal{F}_\beta)} = \text{lf}p(\omega_d^{\text{dual}(P, \emptyset)})|_{\mathcal{F}_\beta}.$$

Since by assumption abduction is not needed, and since the case of co-unfounded sets was handled in Lemma A.1, proving that ABDUAL computes the fixed points specified by these operators is similar to (transfinite) inductions for soundness and completeness of other tabled evaluations of the well-founded semantics. [33] and other papers contain detailed inductions that show that the interpretation induced by the final forest of a tabled evaluation is equivalent to the model preduced by the least fixed point of the operator ω_{ext}^P . Given Theorem 3.1 and Lemma A.1, extending such a proof to abdual is tedious but straightforward. \diamond

The next step is to extend the results of Lemma A.1 to arbitrary abductive frameworks with non-empty sets of abducibles and integrity rules. The results must now prove equivalences to models based on abductive scenarios. In the definition of abductive scenarios (Definition 2.7) a program $P_\mathcal{B}$

is constructed based on a subset of abducible objective literals. A subset \mathcal{B} of the abducibles of an abductive framework σ can be obtained directly from the context, $Context$, of an abductive subgoal; and a program $P_{\mathcal{B}}$ can be generated from \mathcal{B} as in Definition 2.7. In the following lemma, we refer to the program produced by the construction of Definition 2.7 on the objective literals of $Context$ simply as $P_{Context}$, and to $Context$ seen as an interpretation as $\mathcal{I}_{Context}$.

Lemma A.3 *Let $\langle P, \mathcal{A}, Int \rangle$ be an abductive framework.*

1. *Let \mathcal{F} a forest in an ABDUAL evaluation of a query Q against $\langle P, \mathcal{A}, Int \rangle$; S^A be a co-unfounded set of answers in \mathcal{F} ; and let $Context = \bigcup \{C_i \mid \langle S_i, C_i \rangle \text{ :- } DL \mid DL \in S^A\}$. Then there is a minimal unfounded set S^O for $P \cup P_{Context} \cup Int$ with respect to $(\mathcal{I}_{induced(\mathcal{F})} \cup \mathcal{I}_{Context})$ such that $\{S_i \mid \text{not}(S_i) \in \text{heads}(S^A)\} \subseteq S^O$.*
2. *Let $\langle P, \mathcal{A}, \mathcal{B}, Int \rangle$ be an abductive scenario, and \mathcal{I} an interpretation of $P \cup P_{\mathcal{B}} \cup Int$, such that $\mathcal{I}|_{\mathcal{A}}$ is consistent. Let S^O be an unfounded set of objective literals with respect to \mathcal{I} . Let \mathcal{F} be a ABDUAL forest of any query to $\langle P, \mathcal{A}, Int \rangle$ such that*
 - (a) *for each $S_i \in S^O$, there is a tree for $\text{not}(S_i)$ in \mathcal{F} ;*
 - (b) *for all $L \in \mathcal{I}$, if $\langle L, \emptyset \rangle$ is the root of a tree in \mathcal{F} , then $L \in \mathcal{I}_{induced(\mathcal{F})}$; and*
 - (c) *no ABDUAL operations are applicable to \mathcal{F} .*

Then \mathcal{F} will contain an unconditional answer

$$\langle \text{not}(S_i), C_i \rangle \text{ :- } \mid$$

for each $S_i \in S^O$, such that $P_{\bigcup_{i \in S^O} C_i} \subseteq P_{\mathcal{B}}$.

Proof: As in Lemma A.1 we first restrict our attention to finite programs, for all of which the dual program, $udual(P, \mathcal{A})$, formed via Definition A.1 can be used. We then indicate how the arguments can be extended to the dual form of Definition A.1 which is necessary for infinite programs.

1. Let $N_{leaf} = \langle \text{not}(S_i), C_i \rangle \text{ :- } DL \mid$ be an answer in the co-unfounded set of answers S^A . By the construction of unfolded dual transformation (Definition A.1), and by the definitions of ABDUAL it can also be seen that the node $\langle \text{not}(S_i), C_i \rangle \text{ :- } DL \mid$ is a descendant of a node $N = \langle \text{not}(S_i), \rangle \text{ :- } L_1, \dots, L_n$ in which, for each rule r_j with head S_i in $P \cup I$, L_j is a default conjugate of some literal in the body of r_j . (The situation is analogous to that depicted in Figure 8). Now we consider the classes of goal literals in the node N . Those that are in N but not in the *DelayList*, were not resolved away via an ABDUCTION operation, and were not resolved away via an ANSWER CLAUSE RESOLUTION resolution using an answer with a non-abductive context form witnesses of unusability for some rule for S_i in $P \cup I$ by the same argument as in case (1a) in the proof of Lemma A.1. Similarly, those literals in the *DelayList* of N_{leaf} form witnesses of unusability by the same argument as in case (1b) of that proof. This leaves literals that were resolved away via ABDUCTION operations, or were resolved away via an ANSWER CLAUSE RESOLUTION resolution with an answer with a non-abductive context. Either

operation will union the abductive context of a parent node with new objective literals to produce a new child node. We first note that by Definition 3.6, the union of the contexts of all answers in any co-unfounded set is consistent so that the single abductive context, C_i , must also be consistent. Let $O_i \in C_i$ be a given abducible objective literal. In the first case, O_i is the explicit conjugate of some literal in a rule r_j of S_i in $P \cup I$ added to C_i directly through an ABDUCTION operation. In the second case, O_i is necessary to derive an answer that was used via ANSWER CLAUSE RESOLUTION with a coherency axiom or other program or integrity rules to remove a literal from r_j , and so form a witness of unusability. Since abductive contexts must be consistent, O_i is true in $I_{Context}$ (i.e. the interpretation induced by the union of abductive contexts of all answers in S^A), iff it is a witness of unusability for r_i with respect to $\mathcal{I}_{induced} \cup I_{context}$.

2. Again, for the case of finite abductive scenarios, the argument is essentially similar to that of Lemma A.1, but with the addition of abducibles. The only difference is to ensure that the union of contexts of all nodes in the co-unfounded set of answers corresponding to S^O is consistent, which fact follows from Definition 2.7 which implies that the interpretation of abducibles in an abductive scenario gives rise to a consistent interpretation of these abducibles once their truth values are propagated to default literals through coherency.

Extending the proofs to the case of infinite programs means taking account of folding rules created by the use of the dual transformation of Definition 3.1, as discussed in Lemma A.1. It is straightforward to see that the folding rules ensure do not affect consistency of abductive contexts. In addition, since abducibles are propagated through folding rules (and all other rules), and since there can only be a finite number of abducibles to be propagated, extension of the rest of the argument is straightforward.

◇

We next prove a restricted form of Theorem 3.2, which assumes that the final forests exist for an ABDUAL evaluation of a query to these forests. It uses the notion of the model $M(\sigma)$ of abductive solution σ as introduced in Definition 2.8, and of a rule $query :- not(\perp), Q$ as introduced in Definition 3.8.

Lemma A.4 *Let \mathcal{E} be an ABDUAL evaluation of a query Q against an abductive framework $\langle P, \mathcal{A}, Int \rangle$, whose final forest is \mathcal{E}_β . Then $\langle query, Set \rangle :- |$ is an answer in \mathcal{E}_β iff $\sigma = \langle P, \mathcal{A}, Set', I \rangle$ is an abductive solution such that $M(\sigma) \models Q$ and $Set \subseteq Set'$.*

Proof: (Sketch) Given Lemmas A.3 and Lemma A.2 the proof is straightforward. Soundness (\Rightarrow) is shown by an induction on the length of the ABDUAL evaluation, such that each literal in the induced interpretation of the ABDUAL forest is also in the well-founded model. Completeness (\Leftarrow) is shown by a double induction on construction of $M(\sigma)$, the well-founded model of an abductive scenario with the the operators of Definition 3.10 used in the inner induction, and the operator of Definition 3.11 used in the outer induction. ◇

Theorem 3.2 *Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework, and \mathcal{E} be an ABDUAL evaluation of Q against $\langle P, \mathcal{A}, I \rangle$. Then*

1. \mathcal{E} will have a final forest \mathcal{E}_β ;

2. if $\langle \text{query}, \text{Set} \rangle$:- is an answer in \mathcal{E}_β $\sigma = \langle P, \mathcal{A}, \text{Set}, I \rangle$ is an abductive solution for $\langle P, \mathcal{A}, I \rangle$;
3. if $\langle P, \mathcal{A}, \text{Set}, I \rangle$ is a minimal abductive solution for Q , then $\langle \text{query}, \text{Set} \rangle$:- is an answer in \mathcal{E}_β .

Proof: (Sketch)

1. The first statement follows from an argument similar to that made for extended SLG trees in [33]. Briefly recapitulated, it can be seen that all ABDUAL trees are of finite depth, therefore they must have at most a countably infinite number of nodes (e.g. see H. Rogers, Theory of Recursive Functions and Effective Computations, MIT press, 1987 Section 16.3). Furthermore there are at most a countably infinite number of ABDUAL trees in a ABDUAL forest. Now at each successor ordinal of a ABDUAL evaluation, each ABDUAL operation either creates a new tree or adds at least one node to an existing tree. At each limit ordinal the union of all forests indexed by lower ordinals is taken. Therefore, a ABDUAL evaluation can have at most a countably infinite number of states. Thus, the ordinal β is reachable via the transfinite induction of Definition 3.8.
2. Both the second and third statements are immediate from Lemma A.4.

◇

A.2 Proof of Theorems in Section 3.4

Definition A.4 Let P be a program or dual program that is finite and ground. $\text{size}(P)$ denotes the sum, for each rule R in P of 1 plus the number of body literals in R ; $\text{size}(P|_L)$ denotes the size of rules in P whose head is the literal L . $\text{heads}(P)$ denotes the set of literals that occur as heads of rules in P . □

The following lemma shows that the size of the program produced by the dual transformation is linear in the size of the original program. The bound provided is not the tightest possible.

Lemma A.5 Let P be a finite ground extended program and \mathcal{A} a finite set of abducibles. Then $\text{size}(\text{dual}(P \cup \mathcal{A})) < 9\text{size}(P) + 2|\mathcal{A}|$.

Proof: Let O be an objective literal for which there are $m > 0$ rules in P with total size $\text{size}(P|_O)$.

1. Case (1) of Definition 3.1. There will be $m + 2$ rules with $\text{fold}_i^a O$ in their heads or bodies produced by case (1a) of Definition 3.1 for a total size of $3m + 4$. The total size of rules with heads of the form $\text{not}(\text{fold}_i^b O)$ produced by (1b) will be $2(\text{size}(P|_O) - m)$, so that the total number of folding rules for O will be $2\text{size}(P|_O) + m + 4$. Summed up over rules for all objective literals in P , this is $2\text{size}(P) + 4\text{heads}(P) + \text{rules}(P)$, where $\text{rules}(P)$ is the number of rules in P .
2. Case (2). The size of the rules produced by case (2) is bounded by $2(\text{literals}(P) - \text{heads}(P))$.

3. Case (3) Finally, the size of the axioms of coherence is bounded by $|literals(P)| + 2|A|$, where $|literals(P)|$ is the number of literals in P .

Note that $|literals(P)| \leq size(P)$, as is $|heads(P)|$ and $|rules(P)|$, so that $size(dual(P)) < 9size(P) + 2|A|$. \diamond

Lemma A.6 *Let P be a finite ground extended program, and $\langle P, \emptyset, \emptyset \rangle$ be an abductive framework. Let \mathcal{E} be an ABDUAL evaluation of a non-abductive query Q against $dual(P)$, whose final forest is \mathcal{F}_β . Then \mathcal{F}_β can be constructed in at most $2size(dual(P)) + literals(P)$ steps.*

Proof: It takes at most one ABDUAL operation to create a node in \mathcal{F}_β : thus the number of nodes in this forest is an upper bound on the number of ABDUAL operations required to evaluate Q . In \mathcal{F}_β there is at most one tree for each literal in $literals(dual(P))$. The root node N_L for a literal L in $dual(P, \emptyset)$ has one child for each rule for L in P . Consider a child N_R of N_L formed by PROGRAM CLAUSE RESOLUTION using a rule R . Then there are at most $2L_R$ descendants of N_R in \mathcal{F}_β where L_R are the number of body literals in R . To see this, first note that the number of goal literals in N_R is L_R . Further, since $dual(P, \emptyset)$ is ground and the set of abducibles is empty, any descendant N_D of N_R can have at most one child for each of the operations in Definition 3.9. Also consider that any child N_{child} of N_D

1. has the same number of goal literals as N_D and fewer delay literals; (e.g. if N_{child} was produced by SIMPLIFICATION or CO-UNFOUNDED SET REMOVAL); or
2. has the same number of delay literals as N_D and one fewer goal literal (e.g. if N_{child} was produced by ANSWER CLAUSE RESOLUTION); or
3. has one more delay literal than N_D and one fewer goal literal than N_D (e.g. if N_{child} was produced by DELAYING); or
4. is a failure node (e.g. if N_{child} was produced by SIMPLIFICATION).

In cases (1) and (2) above the size of N_D is reduced; in case 3 we note that by Definition 3.9 a literal can be delayed only once in any path from a root of a tree to a leaf. In case 4, we note that no operations are applicable to a failure node. Thus the length of the path from N_O to the unique leaf descendant of N_R in \mathcal{F}_β is at most $2L_R$, leading to the bound in the statement. \diamond

Theorem 3.3 *Let $\langle P, \mathcal{A}, I \rangle$ be an abductive framework such that P and I are finite ground extended programs, and \mathcal{A} is a finite set of abducibles. Let \mathcal{E} be an ABDUAL evaluation of a query Q against $\langle P, \mathcal{A}, I \rangle$. Then \mathcal{E} will have a final forest that is produced by a finite number of ABDUAL operations.*

Proof: By Lemma A.5, $dual((P \cup I), \mathcal{A})$ is finite if $\langle P, \mathcal{A}, I \rangle$ is finite, so we need only consider the direct evaluation of $dual((P \cup I), \mathcal{A})$. In the case in which \mathcal{A} is empty, the statement is immediate from Lemma A.6. In the case in which $|\mathcal{A}| = n$, $n > 0$, then the ABDUCTION operation must be taken into consideration. As in the proof of Lemma A.6, the total number of trees in any forest of \mathcal{E} is bounded by $literals(dual((P \cup I), \mathcal{A}))$. Let N_L be the root node of a tree for L in a forest \mathcal{F} of \mathcal{E} , and

N_R be a child of N_L produced by PROGRAM CLAUSE RESOLUTION using a rule R . As in the proof of Lemma A.6, path to any descendant of N_R is at most $2L_R$, where L_R is the number of body literals in R . This is immediate Lemma A.6 since the ABDUCTION operation reduces the number of goal literals in any node by 1. However, unlike the case in Lemma A.6 where the set of abducibles is empty, the number of children of L_R or its descendants can be greater than one, depending on the number of abductive solutions for a goal literal or set of delay literals. However, the number of these solutions is always finite, so that \mathcal{F} will have a finite number of nodes. As before, each node is produced by a single ABDUAL operation, so that \mathcal{F} can be produced by a finite number of operations. Since an arbitrary forest \mathcal{F} must be finite, a final forest must be produced by a finite number of ABDUAL operations. \diamond

Theorem A.7 *Let \mathcal{F} be the final forest in a ABDUAL evaluation \mathcal{E} of a query Q against an abductive framework $\langle P, \mathcal{A}, I \rangle$. Let $C_{context}$ be the maximal cardinality of the context of any abductive subgoal in \mathcal{F} , and $C_{abducibles}$ be the cardinality of \mathcal{A} . Then \mathcal{F} can be constructed in at most $M \times 2 \text{size}(\text{dual}((P \cup I), \mathcal{A}))$ steps, where*

$$M = \sum_{i \leq C_{context}} \binom{C_{abducibles}}{i}$$

Proof: Let N_L be a root node of a tree for a literal L in \mathcal{F} . As noted in the proof for Lemma A.6, there is one child for N_L for every rule R for L in $\text{dual}((P \cup I), \mathcal{A})$. Such a node, N_R has a number of leaf descendants that is bounded by the number of abductive solutions that are possible to return, via ANSWER CLAUSE RESOLUTION, to the descendants of N_R . By construction, this bound is M . Again by the considerations in Lemma A.6, the length of the path from N_R to any leaf is at most $2R_L$ where R_L is the number of body literals in the rule that produced N_R . Further, the length of the path from N_L to any leaf is $2R_L + 1$, which is 2 times the size of R minus 1. so that the number of nodes in the tree rooted by N_L is bounded by $2M \text{size}(P|_L)$. Summing this for all literals in $\text{dual}((P \cup I), \mathcal{A})$, the bound of the statement is obtained. \diamond