

Dynamic Consistency Checking in Goal-Directed Answer Set Programming

KYLE MARPLE and GOPAL GUPTA

*Department of Computer Science
The University of Texas at Dallas*

submitted 1 January 2003; revised 1 January 2003; accepted 1 January 2003

Abstract

In answer set programming, inconsistencies arise when the constraints placed on a program become unsatisfiable. In this paper, we introduce a technique for *dynamic consistency checking* for our goal-directed method for computing answer sets, under which only those constraints deemed relevant to the partial answer set are tested, allowing inconsistent knowledgebases to be successfully queried. However, the algorithm guarantees that, if a program has at least one consistent answer set, any partial answer set returned will be a subset of some consistent answer set.

KEYWORDS: dynamic consistency checking, answer set programming, goal-directed, consistent query answering

1 Introduction

Answer Set Programming (ASP) (Gelfond and Lifschitz 1988) has gained popularity as a way to develop non-monotonic reasoning applications. Three problems which prevent ASP from being adopted on a larger scale are (i) the need to compute a complete answer set regardless of the query, (ii) the ability of a minor inconsistency to render an entire knowledgebase useless, and (iii) the need to ground programs prior to execution. Our previous work with goal-directed ASP addresses the first (Marple et al. 2012), and we leave the third for future work. In this paper, we address the second problem in context of goal-directed execution of answer set programs.

Currently, most popular ASP solvers rely on SAT solvers (Giunchiglia et al. 2004; Gebser et al. 2007) which can't simply disregard inconsistencies that are unrelated to a query. Because complete answer sets are computed, the underlying program must be consistent. Thus much of the existing work in querying inconsistent knowledgebases has focused on repairing programs to restore consistency (Arenas et al. 2003). In contrast, our goal in this paper is to be able to work with the consistent part of the knowledgebase, i.e., as long as a query does not invoke clauses from the part of the knowledgebase that is inconsistent, we should be able to execute it and produce an answer set, if one exists. Thus, we do deviate from standard ASP semantics, as under ASP semantics, there are no answer sets in the presence of inconsistencies in the knowledgebase.

In this paper, we introduce *dynamic consistency checking* (DCC), a method for querying inconsistent databases that requires no modification of the underlying programs or

queries. Instead, DCC takes advantage of goal-directed answer set programming to ignore inconsistencies that are unrelated to the current query. Additionally, because DCC reduces the number of consistency checks that a partial answer set must satisfy, it can significantly improve the performance of goal-directed execution.

At the core of the problem is the issue of relevance. Because ASP and the underlying stable model semantics lack a *relevance property*, the truth value of an atom can depend on other, totally unrelated rules and atoms (Dix 1995). Because such rules may not be encountered during normal top-down execution, any goal-directed execution strategy for ASP must either alter the semantics or employ some form of consistency checking to ensure correctness. In designing our goal-directed method we chose the latter route, employing consistency checks to ensure that constraints imposed by these rules are satisfied.

DCC employs splitting sets (Lifschitz and Turner 1994) to reduce the number of consistency checks that must be satisfied while retaining strong guarantees regarding correctness. Execution using DCC employs a modified relevance criteria to determine which consistency checks are relevant to the current partial answer set, and only those checks are enforced.

DCC has been implemented as an extension of the *Galliwasp* system (Marple and Gupta 2013), which makes use of our original goal-directed method. As we will demonstrate, DCC has several advantages over other potential strategies based on ignoring unrelated inconsistencies. We will show that, if a program has at least one consistent answer set, then a query will succeed using DCC if and only if the partial answer set returned is a subset of some consistent answer set. If no consistent answer set exists, then DCC can allow partial answer sets to be found for a consistent subset of the program. We will also demonstrate that DCC can improve the performance of goal-directed execution and that partial answer sets produced using DCC can provide more targeted results than either full answer sets or partial answer sets with comprehensive consistency checking.

The remainder of the paper is structured as follows. In Section 2 we discuss issues that are potential impediments to widespread adoption of ASP. Next, in Section 3, we give an overview of goal-directed ASP, focusing on consistency checking. In Section 4 we introduce our technique for dynamic consistency checking using splitting sets and prove several interesting properties. In Section 5 we examine advantages of DCC and compare the results of *Galliwasp* with and without dynamic consistency checking. Finally, in Section 7 we discuss related and future work and draw conclusions.

2 Answer Set Programming: Challenges

While the Answer Set Programming paradigm has gained wide popularity among researchers, there are still issues that stand in the way of its use by ordinary users. The overarching goal of our research project is to eliminate such issues. The major issues are briefly described next, though this paper is mainly concerned with addressing only the last one.

The first problem relates to grounding an ASP program. Because existing systems for executing ASP programs rely on SAT solvers, ASP programs containing predicates have to be grounded first. Even when restricted to finitely groundable programs, the size of a grounded program can be exponentially large. Thus, while writing ASP programs, one has to write code in a way that will keep the size of the grounded program small. The

grounding step can be avoided if goal-directed strategies, such as Galliwasp, are developed and used to execute ASP programs. At present, even though the execution algorithm used by Galliwasp is goal-directed, it assumes that the input program is grounded. Note that work is in progress to extend Galliwasp so that predicate ASP programs (including those containing functions) can be executed in a goal-directed manner without being grounded (Salazar et al. 2014).

The second problem relates to computing an entire model of a program. Most current ASP execution methods compute the *entire* answer set, but in practice, we may only be interested in knowing if a specific piece of knowledge can be inferred. Consider the case of a large relational database coded in ASP. Without additional constraints, a complete answer set will contain all of the information in the database, not just the answer to a successful query. To work around this, constraints will need to be added to pare down the results, effectively requiring that a program be written where a single query might otherwise suffice. So if we rely on such solvers, then ASP can be used for solving specific problems, but its use for building large knowledge-based applications will pose challenges.

Finally, the third problem relates to being able to work with ASP programs which are inconsistent. As long as the answer being sought only depends on a consistent subset of the knowledgebase, one should be able to infer that knowledge. However, this is not the case with current ASP systems. The entire knowledgebase has to be consistent in order for them to produce a solution. To take a trivial example, consider a consistent ASP program κ to which the clause $p :- \text{not } p.$ is added, where p does not occur elsewhere in the program. The augmented program will have no answer sets. It is difficult for SAT solver-based approaches to identify subsets of the program that are consistent. A query-driven, goal-directed approach, in contrast, only ‘touches’ those parts of the program that are needed for establishing the query. All constraints that involve any of the literals ‘touched’ during the execution of the query, directly or indirectly, must also be enforced. However, constraints that do not involve such literals need not be executed, as they are independent of the part of the program that was involved in answering the query. This is precisely the idea behind our work on *dynamic consistency checking* presented in this paper: only consistency checks that involve the portion of the program that is ‘touched’ by the query are executed. Thus, adding the rule $p :- \text{not } p.$ to knowledgebase κ above will not alter the execution of the program unless a query contains p .

3 Goal-Directed Answer Set Programming

Under our basic goal-directed method, a partial answer set is constructed by adding both positive and negative literals as they succeed during execution. When a query succeeds and all consistency checks have been satisfied, the set of positive literals in the partial answer set is guaranteed to be a subset of some consistent answer set of the program (Marple et al. 2012). This method can be applied to arbitrary ASP programs, including those with rules that contain classical negation and disjunction. Such rules are simply converted to an equivalent set of normal rules.

Execution uses a modified form of co-*SLD* resolution (*SLD* resolution with coinduction) (Gupta et al. 2007). Under co-*SLD* resolution, each call is added to the *coinductive hypothesis set* (CHS); a call can succeed *coinductively* if it unifies with an ancestor call in the CHS. In our goal-directed execution method, the CHS also serves as the candidate

answer set. However, some modifications are necessary to adapt co-SLD resolution to ASP:

- Negated calls are also allowed to succeed coinductively, i.e., negated calls (e.g., not p) are added to the CHS. A negated call can succeed coinductively if it unifies with an ancestor negated call in the CHS.
- A literal and its negation cannot be in the CHS at the same time. If adding a literal to the CHS leads to such a situation, the computation fails and backtracking ensues.
- Coinductive success is allowed only if an even, non-zero number of negations occur between the recursive call and its ancestor call.

While the above description covers the basic execution of our algorithm, it omits perhaps the most important part, consistency checking. To understand the role of consistency checking, we must first examine the issue of relevance in more detail.

3.1 Relevance

The issue of relevance is central to goal-directed ASP. In defining relevance, (Dix 1995) uses the dependency graph of a program P and the following notions:

- “ $dependencies_of(X) := \{A : X \text{ depends on } A\}$ ”, i.e. X calls A directly or indirectly, and
- “ $rel_rule(P, X)$ is the set of *relevant rules* of P with respect to X , i.e. the set of rules that contain an $A \in dependencies_of(X)$ in their heads.”

Then, “given any semantics SEM and a program P , it is perfectly reasonable that the truth-value of a literal L , with respect to $SEM(P)$, only depends on the subprogram formed from the relevant rules of P with respect to L ”, formalized as:

Definition 1

“Relevance states that for all literals L : $SEM(P)(L) = SEM(rel_rule(P, L))(L)$.” (Dix 1995)

Despite being “perfectly reasonable”, the above definition of relevance does not hold for ASP. This is due to the presence of rules which contain an odd loop over negation (OLON). OLONs occur implicitly in rules with an empty head, but also occur in rules with non-empty heads, whenever a rule can be called recursively with an odd number of negations between the original and recursive calls. These “OLON rules” place constraints on a program that must be satisfied by any consistent answer set. For example, given an OLON rule of the form:

$p :- B, \text{ not } p.$

where B is a conjunction of literals, one of the following must be satisfied:

1. p must succeed through other means, or
2. at least one literal in B must fail.

That is, the rule imposes the constraint $p \vee \text{not } B$ on the program. Such a rule can thus alter the truth-value of a literal in B despite not being relevant to the literal under Definition 1.

```

p :- q.                % Rule 1: OLON
q :- not r, not p.     % Rule 2: OLON
r :- not p.            % Rule 3: Ordinary
:- q, r.               % Rule 4: OLON

chk_1 :- p.
chk_1 :- not q.
chk_2 :- r.
chk_2 :- p.
chk_2 :- q.
chk_4 :- not q.
chk_4 :- not r.
nmr_check :- chk_1, chk_2, chk_4.

```

Fig. 1. A simple program with consistency checks added.

3.2 Consistency Checking

Because ASP lacks a relevance property of its own, our algorithm uses consistency checks to enforce a modified relevance property, where for a program P and literal L , the set of rules in P relevant to L is expanded to include every OLON rule in the program (Marple et al. 2012). That is,

$$nmr_rel_rul(P, L) = rel_rul(P, L) \cup OLON(P) \quad (1)$$

where $rel_rul(P, L)$ is the set of relevant rules defined in Section 3.1 and $OLON(P)$ is the set of OLON rules in P . The semantics of P with respect to L can now be defined in terms of the subprogram formed by the expanded set of relevant rules:

$$SEM(P)(L) = SEM(nmr_rel_rule(P, L))(L) \quad (2)$$

This property ensures that an answer set of the subprogram, if one exists, will be a subset of some consistent answer set of P (Marple et al. 2012).

To enforce our modified relevance property, our method uses a special rule, the *non-monotonic reasoning check* (NMR check), which calls a sub-check for each OLON rule in a program. Each sub-check ensures that the associated OLON rule is satisfied. The NMR check is then automatically appended to each query, ensuring that the property will hold for any query which succeeds (Marple et al. 2012).

The construction of the sub-checks involves creating rules for the dual of each OLON rule in the program. Duals explicitly encode the negation of a literal. For example, given:

```
p :- q, not r.
```

the dual rules for p are:

```
not p :- not q.
not p :- r.
```

In the case of sub-checks, the negation of the head is first appended to encapsulate the success of a literal through other means. The entire process is as follows:

1. For rules with non-empty heads, the negation of the head is appended to the body of the rule, if not already present.
2. The dual of the rule is computed.
3. The dual is given a unique head, which is also added to the body of the NMR sub-check.

The example in Figure 1 shows a simple ASP program with the NMR check and sub-checks added.

While this ensures that our method adheres to the semantics of ASP, the execution of the NMR check can adversely impact performance. ASP programs routinely make heavy use of headless rules to enforce constraints, which can result in an NMR check which contains thousands of goals. For example, an instance of the 20-Queens problem can produce an NMR check containing 25,100 goals, each representing a sub-check that must be executed alongside any query. A means of reducing the performance impact of these checks is thus extremely desirable.

4 Dynamic Consistency Checking

Dynamic Consistency Checking (DCC) began as an attempt to improve the performance of goal-directed execution. While we have developed various other techniques to reduce the performance impact of consistency checking, none of them reduce the actual number of checks that must be satisfied, as this is impossible to do while guaranteeing full compliance with the ASP semantics. DCC was our attempt to reduce the number of checks performed while staying as close to the original ASP semantics as possible.

As any reduction in the number of consistency checks will result in non-compliance with the ASP semantics, selecting which checks to enforce depends on the properties desired from the modified semantics. In the case of DCC, these properties also make the technique useful for querying inconsistent knowledgebases.

Definition 2

For a program P , the *desired properties* of DCC are:

1. Execution shall always be consistent with the ASP semantics of the sub-program of P (further defined in Section 4.1).
2. If P has at least one consistent answer set, execution shall be consistent with the ASP semantics of P .

In this section, we discuss the relevance property employed by DCC before moving on to the algorithm itself. Finally, we provide proofs that DCC satisfies the above properties.

4.1 Relevance Under DCC

While our original relevance property, given in Formula 2, makes every consistency check relevant to every literal, DCC selects only those checks necessary to enforce our desired properties from Definition 2. Relevant checks are dynamically selected based on the literals in the partial answer set.

At first glance, it might seem sufficient to select only those checks which directly call literals in the partial answer set (or their negations). However, this can lead to

```

a :- b.
b :- not c.
c :- not b.

p :- a.
q :- b.
:- p, q.

chk_1 :- not p.
chk_1 :- not q.
nmr_check :- chk_1.

```

Fig. 2. Example program (consistency checks added).

```

:- p, q.
q :- not r, not q.

chk_1 :- not p.
chk_1 :- not q.
chk_2 :- r.
chk_2 :- q.
nmr_check :- chk_1, chk_2.

```

Fig. 3. Example program (consistency checks added).

incorrect results. Consider the program in Figure 2. One consistent answer set exists: $\{c, \text{not } a, \text{not } b, \text{not } p, \text{not } q\}$. However, given a query $?- a.$, selecting only those checks which directly call some literal in the partial answer set will yield $\{a, b, \text{not } c\}$, thus violating our desired properties.

Clearly, our properties require that we select at least those checks which can potentially reach a literal in the partial answer set. However, this can lead to behavior that is difficult to predict. Consider the program in Figure 3 with the query $?- \text{not } p.$ The presence of either q or $\text{not } q$ in each OLON rule might seem to indicate that both consistency checks will be activated and cause the query to fail. However, only `chk_1` will be activated. Because the first clause will succeed, neither q nor its negation will be added to the partial answer set, and the query will succeed.

To achieve more predictable behavior, DCC selects relevant checks using specially constructed *splitting sets*. A splitting set for a program is any set of literals such that if the head of a rule is in the set, then every literal in the body of the rule must also be in the set (Lifschitz and Turner 1994). The rules in a program P can then be divided relative to a splitting set U into the bottom, $b_U(P)$, containing those rules whose head is in U , and the top, $P \setminus b_U(P)$.

The splitting sets used to determine relevant NMR sub-checks are created by constructing splitting sets for each NMR sub-check and merging sets whose intersection is non-empty. The result is a set of disjoint splitting sets U_i such that for an NMR sub-check

C , if $C \in U_i$, then for every literal L reachable by C , $L \in U_i$. This allows us to define the sub-checks relevant to a literal as those whose heads are in the same splitting set:

$$\begin{aligned} dcc_rel_rul(P, L) &= rel_rul(P, L) \cup OLON(P, L), \\ OLON(P, L) &= \{R : R \in OLON(P) \cap b_{U_i}(P) \wedge L \in U_i\} \end{aligned} \quad (3)$$

where $OLON(P, L)$ is the set of OLON rules relevant to L . This leads us to DCC's relevance property, which defines the semantics of P with respect to L in terms of the new set of relevant rules:

$$SEM(P)(L) = SEM(dcc_rel_rule(P, L))(L) \quad (4)$$

This definition allows for more predictable behavior than simply selecting the checks reachable by a given literal. In the case of Programs 2 and 3, only one splitting set will be created, resulting in behavior that is identical to normal goal-directed ASP. Indeed, as we will prove in Section 4.3, execution will be consistent with ASP whenever a program has at least one answer set.

4.2 Execution with DCC

Given DCC's relevance property in Formula 4, our goal-directed execution strategy must be modified to enforce it. A query should succeed if and only if every OLON rule relevant to a literal in the partial answer set is satisfied. In addition to creating the associated splitting sets, the application of the relevant NMR sub-checks also becomes more complex.

The creation of the necessary splitting sets can be accomplished by examining a program's call graph after the NMR sub-checks have been added. A simple depth-first search is sufficient to construct the splitting set for an individual sub-check, after which overlapping sets can be merged. For added efficiency, constructing and merging the sets can be performed simultaneously: whenever a literal is encountered that has already been added to another set, that set is merged with the current one. This eliminates the need to traverse any branch in the call graph more than once. The overhead of searching the sets themselves can be minimized with proper indexing.

To apply the NMR check when executing a query with DCC, it must also be dynamically constructed. The NMR check should consist of those sub-checks which are relevant to a literal in the partial answer set. However, because the sub-checks themselves may add literals to the partial answer set, simply executing the query and then selecting the relevant checks once is insufficient. Instead, each time a literal succeeds, the relevant sub-checks are added to the NMR check. Similarly, the state of the NMR check is restored when backtracking occurs. In this manner, the NMR check will always remain consistent with the current partial answer set.

4.3 Correctness of DCC

Now that we have established DCC's algorithm, we can prove that it satisfies the property it was designed to enforce. That is:

Theorem 1

If a program P has at least one consistent answer set, then a query will succeed under DCC if and only if the partial answer set is a subset of some consistent answer set of P .

Proof

Observe that, if a DCC query succeeds, the partial answer set will be $X = A \cup B$ where

- A is a partial answer set of the splitting set U formed by the union of the splitting sets containing relevant NMR sub-checks
- B is the set of succeeding literals which are not reachable by any NMR sub-check

Per the Splitting Set Theorem (Lifschitz and Turner 1994), a set X' is an answer set of P if and only if $X' = A' \cup B'$ where A' is an answer set of $b_U(P)$, B' is an answer set of $e_U(P \setminus b_U(P), A')$, and $A' \cup B'$ is consistent.¹ Thus our theory will hold if $A \subseteq A'$, $B \subseteq B'$ and $A' \cup B'$ is consistent.

Because every NMR sub-check relevant to some literal in A will be activated and must succeed for the DCC query to succeed, A will always be a subset of some consistent answer set of $b_U(P)$. Furthermore, such an answer set must exist for the DCC query to succeed. Thus, for any succeeding DCC query, there exists an answer set A' of $b_U(P)$ such that $A \subseteq A'$.²

Because only OLON rules can lead to inconsistency in an ASP program³, the set B will always be a subset of some consistent answer set of $e_U(P \setminus b_U(P), A')$, if one exists. Therefore, if at least one consistent answer set exists for P , we can select B' such that B' is an answer set of $e_U(P \setminus b_U(P), A)$ such that $B \subseteq B'$.

Finally, because A' contains every NMR sub-check relevant to any literal in A , A' will always be consistent with B' . Thus, if P has at least one answer set, a query will succeed under DCC if and only the partial answer set is a subset of some consistent answer set of P . $\square$

5 Advantages of DCC

Execution with DCC offers several advantages over normal goal-directed ASP. The three primary advantages are partial answer sets of inconsistent programs, output that is relevant to the query, and improved performance.

5.1 Answer Sets of Inconsistent Programs

One disadvantage of ASP is the way in which it handles inconsistency in a knowledgebase. Any inconsistency, no matter how small, renders the entire program inconsistent, and thus no answer set will exist. This behavior can be particularly inconvenient in large knowledgebases where an inconsistency may be completely unrelated to a particular query. Given a large, perfectly consistent database implemented in ASP, adding the rule $:- \text{not } c.$ where c is a unique literal, will cause any query to the database to fail.

With DCC, if a query succeeds prior to adding the rule above, then it will continue to succeed even after the rule is added.

¹ For a set X of positive literals in U , $e_U(P \setminus b_U(P), X)$ is a partial evaluation of the top of P with respect to X . The partial evaluation is constructed by first dropping rules whose bodies contain the negation of a literal in X and then removing calls to literals in X from the bodies of the remaining rules.

² If no literals in the query are reachable by any NMR sub-checks, U will be empty and both A' and A will be the empty set.

³ While rules involving classical negation and disjunction can lead to inconsistency, Galliwasp handles these by converting them to a set of equivalent normal rules, including OLON rules.

Table 1. *Comparative Performance Results*

Problem	Splitting Sets	Query	Execution Times ^a	
			Original	w/ DCC
hanoi-5x15	0	solveh	0.276	0.274
pigeons-30x30	1	solvep	0.065	0.065
schur-3x13	1	solves	0.105	0.105
hanoi-schur	1	solveh	0.134	0.028
hanoi-schur	1	solves	0.134	0.134
hanoi-pigeons	1	solveh	0.346	0.341
hanoi-pigeons	1	solvep	0.343	0.342
pigeons-schur	2	solvep	9.958	0.672
pigeons-schur	2	solves	9.745	0.172
han-sch-pigs	2	solveh	9.817	0.093
han-sch-pigs	2	solvep	9.780	0.094
han-sch-pigs	2	solves	9.942	0.201

^a CPU time in seconds.

5.2 Query-relevant Output

One advantage of goal-directed ASP is the ability to compute partial answer sets using a query. Ideally, partial answer sets will contain only literals which are related to the query. However, the execution of the NMR check can force the addition of literals which are unrelated to the current query. By omitting unnecessary NMR checks, DCC can limit this irrelevant output.

Consider the case where two consistent ASP programs, A and B, are concatenated to form a new program C. Assume that A and B have no literals in common and that each contains one or more OLON rules. A full answer set of C will obviously contain literals from both of the sub-programs. As a result of the OLON rules, any partial answer set obtained using goal-directed ASP will also contain literals from both sub-programs. However, using DCC, a succeeding query which targets only one sub-program will only contain literals from that sub-program.

Exploiting this behavior does require care on the part of the programmer. For example, many ASP programs use OLON rules in place of queries. However, such rules will often force all or most of a program's literals into a single splitting set. As a result, every OLON rule will always be deemed relevant, and DCC will function no differently than normal goal-directed ASP. We will see this behavior in some of the sub-programs examined in the next section.

5.3 Performance Compared to Normal Consistency Checking

In this section we compare *Galliwasp*'s performance on several programs, with and without DCC. As the results in Table 1 demonstrate, programs that take advantage of DCC can see a massive improvement in performance. Additionally, even when a program does not take advantage of DCC, the overhead remains minimal.

To simulate programs which take advantage of DCC, the following three programs were concatenated together in various combinations:

- **hanoi-5x15** is a 5 ring, 15 move instance of the Towers of Hanoi. The query `?- solveh.` will return a partial answer set containing the solution.
- **pigeons-30x30** is an instance of the MxN-Pigeons problem. The query `?- solvep.` will find a complete answer set.
- **schur-3x13** is a 3 partition, 13 number instance of the Schur Numbers problem. The query `?- solves.` finds a complete answer set.

Each of the three base programs, and thus each combination, has at least one consistent answer set. The Towers of Hanoi instance contains no OLON rules, and consequently no splitting sets. The other two programs contain OLON rules that force the computation of a complete answer set, and thus have one splitting set each. As a result, a DCC query containing only `solveh` will not activate any NMR sub-checks, while queries containing `solvep` or `solves` will activate every NMR sub-check for their respective problems. Thus DCC execution of `solveh` will not access any splitting sets, while `solvep` and `solves` will access one set each.

In general, the fewer splitting sets accessed by a DCC query relative to the total, the better it will perform compared to a non-DCC query. This is exemplified by the cases with two splitting sets in Table 1. In the programs tested, each splitting set represents a large number of OLON rules. As the non-DCC results indicate, the negative impact of increasing the number of OLON rules can be immense. DCC is able to avoid this by satisfying only those rules relevant to the current query.

6 Related and Future Work

DCC is an extension of goal-directed ASP (Marple et al. 2012) and has been implemented using the *Galliwasp* system (Marple and Gupta 2013). The technique relies heavily on the properties of splitting sets, and the Splitting Set Theorem in particular (Lifschitz and Turner 1994).

Numerous other methods for querying inconsistent databases have been developed. The problem of Consistent Query Answering is defined in terms of minimal database repairs in (Arenas et al. 1999), which develops a technique based on query modification that is built upon in several subsequent works (Celle and Bertossi 2000; Arenas et al. 2003). However, these techniques require that database inconsistencies be identified and accounted for. Because DCC relies on a goal-directed technique for computing answer sets, our method allows inconsistent information to simply be ignored unless it directly relates to the current query.

Plans for future work focus on modifying the technique to work with ungrounded ASP programs. Detecting OLONs and constructing the associated splitting sets prior to grounding has the potential to both reduce the overhead and allow the use of DCC with a wider range of solvers. Of particular interest is integration with a datalog ASP system currently under development (Salazar et al. 2014).

7 Conclusions

In this paper we have introduced Dynamic Consistency Checking (DCC), a technique for querying inconsistent ASP programs using a goal-directed execution method. We have discussed the relevant aspects of goal-directed ASP, presented the relevance criteria which DCC enforces and proven that DCC is consistent with the ASP semantics for programs which have at least one consistent answer set. Additionally, we have examined the advantages of DCC with respect to querying inconsistent databases, achieving more useful output from queries, and improving the performance of the *Galliwasp* system. As our results demonstrate, DCC can be efficiently implemented and programs which take advantage of it can achieve significant benefits. Future work will focus on allowing DCC to operate on ungrounded ASP programs and adapting the technique into additional ASP solvers.

References

- ARENAS, M., BERTOSSI, L., AND CHOMICKI, J. 1999. Consistent Query Answers in Inconsistent Databases. In *Proceedings of the Eighteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. PODS '99. ACM, New York, NY, USA, 68–79.
- ARENAS, M., BERTOSSI, L., AND CHOMICKI, J. 2003. Answer Sets for Consistent Query Answering in Inconsistent Databases. *Theory Pract. Log. Program.* 3, 4 (July), 393–424.
- CELLE, A. AND BERTOSSI, L. E. 2000. Querying Inconsistent Databases: Algorithms and Implementation. In *Proceedings of the First International Conference on Computational Logic*. CL '00. Springer-Verlag, London, UK, UK, 942–956.
- DIX, J. 1995. A Classification Theory of Semantics of Normal Logic Programs: II. Weak Properties. *Fundamenta Informaticae* 22, 257–288.
- GEBSER, M., KAUFMANN, B., NEUMANN, A., AND SCHAUB, T. 2007. Clasp: A Conflict-Driven Answer Set Solver. In *Proceedings of the 9th international conference on Logic Programming and Nonmonotonic Reasoning*. LPNMR'07. Springer-Verlag, 260–265.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The Stable Model Semantics for Logic Programming. In *Proceedings of the Fifth international conference on Logic Programming*. MIT Press, 1070–1080.
- GIUNCHIGLIA, E., LIERLER, Y., AND MARATEA, M. 2004. SAT-Based Answer Set Programming. In *Proceedings of the 19th national conference on Artificial Intelligence*. AAAI'04. AAAI Press, 61–66.
- GUPTA, G., BANSAL, A., MIN, R., SIMON, L., AND MALLYA, A. 2007. Coinductive Logic Programming and Its Applications. In *Proceedings of the 23rd international conference on Logic Programming*. ICLP'07. Springer-Verlag, 27–44.
- LIFSCHITZ, V. AND TURNER, H. 1994. Splitting a Logic Program. In *Proceedings of the eleventh international conference on Logic programming*. ICLP '94. MIT Press, Cambridge, MA, USA, 23–37.
- MARPLE, K., BANSAL, A., MIN, R., AND GUPTA, G. 2012. Goal-directed Execution of Answer Set Programs. In *Proceedings of the 14th symposium on Principles and practice of declarative programming*. PPDP '12. ACM, New York, NY, USA, 35–44.
- MARPLE, K. AND GUPTA, G. 2013. Galliwasp: A Goal-Directed Answer Set Solver. In *Logic-Based Program Synthesis and Transformation*. Lecture Notes in Computer Science, vol. 7844. Springer Berlin Heidelberg, 122–136.
- SALAZAR, E., MARPLE, K., AND GUPTA, G. 2014. Galliwasp II: Goal-directed Execution of Predicate Answer Set Programs. Tech. rep., The University of Texas at Dallas. Forthcoming.