

Grounding and Solving in Answer Set Programming

Benjamin Kaufmann, Nicola Leone, Simona Perri, Torsten Schaub

■ *Answer set programming is a declarative problem-solving paradigm that rests upon a work flow involving modeling, grounding, and solving. While the former is described by Gebser and Schaub (2016), we focus here on key issues in grounding, or how to systematically replace object variables by ground terms in an effective way, and solving, or how to compute the answer sets, of a propositional logic program obtained by grounding.*

Answer set programming (ASP) combines a high-level modeling language with effective grounding and solving technology. Moreover, ASP is highly versatile by offering various elaborate language constructs and a whole spectrum of reasoning modes. The work flow of ASP is illustrated in figure 1.

At first, a problem is expressed as a logic program. A grounder systematically replaces all variables in the program by (variable-free) terms, and the solver takes the resulting propositional program and computes its answer sets (or aggregations of them).

ASP's success is largely due to the availability of a rich modeling language (Gebser and Schaub 2016) along with effective systems. Early ASP solvers SModels (Simons, Niemelä, and Soinen 2002) and DLV (Leone et al. 2006) were followed by SAT¹-based ones, such as ASSAT (Lin and Zhao 2004) and Cmodels (Giunchiglia, Lierler, and Maratea 2006), before genuine conflict-driven ASP solvers such as clasp (Gebser, Kaufmann, and Schaub 2012a) and WASP (Alviano et al. 2015) emerged. In addition, there is a continued interest in mapping ASP onto solving technology in neighboring fields, like SAT or even MIP² (Janhunen, Niemelä, and Sevalnev 2009; Liu, Janhunen, and Niemelä 2012), and in the automatic selection of the appropriate solver by heuristics (Maratea, Pulina, and Ricca 2014).

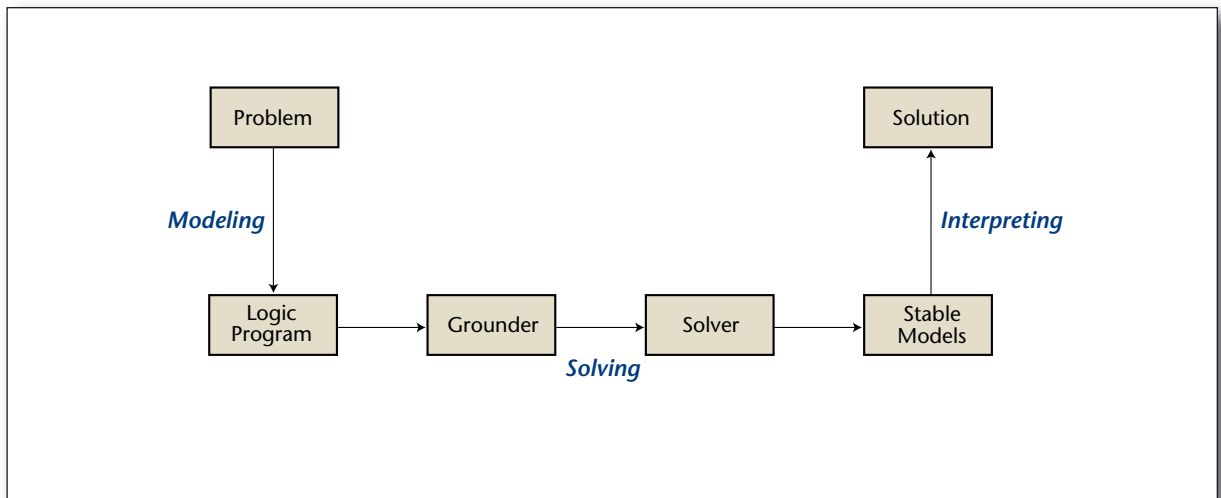


Figure 1. The Work Flow of Answer Set Programming.

By contrast, modern grounders like (the one in) DLV (Faber, Leone, and Perri 2012) or GrinGo (Gebser et al. 2011) are based on seminaïve database evaluation techniques (Ullman 1988) for avoiding duplicate work during grounding. Grounding is seen as an iterative bottom-up process guided by the successive expansion of a program’s term base, that is, the set of variable-free terms constructible from the signature of the program at hand. Other grounding approaches are pursued in GIDL (Wittocx, Mariën, and Denecker 2010), Lparse (Syrjänen 2001), and earlier versions of GrinGo (Gebser, Schaub, and Thiele 2007). The latter two bind nonglobal variables by domain predicates to enforce ω - or λ -restricted (Syrjänen 2001; Gebser, Schaub, and Thiele 2007) programs that guarantee a finite grounding, respectively.

In what follows, we describe the basic ideas and major issues of modern ASP grounders and solvers, also in view of supporting ASP’s language constructs and reasoning modes.

Grounding

Modern ASP systems perform their computation by first generating a ground program that does not contain any variable but has the same answer sets as the original program. This phase, usually referred to as *grounding* or *instantiation*, solves a complex problem. In the case in which input nonground programs can be assumed to be fixed (data complexity), this task is polynomial. However, as soon as variable programs are given in input, grounding becomes EXPTIME-hard, and the produced ground program is potentially of exponential size with respect to the input program. To give an idea of that, consider the following program containing only one rule, and two facts:

```

obj(0). obj(1).
tuple(X1, ..., Xn) :- obj(X1), ..., obj(Xn).
  
```

The ground instantiation of the rule contains 2^n ground rules, corresponding to the number of n -tuples, over a set of two elements. For more details about complexity of ASP the reader may refer to Dantsin et al. (2001).

Grounding, hence, may be computationally very expensive having a big impact on the performance of the whole system, as its output is the input for an ASP solver, that, in the worst case, takes exponential time in the size of the input. Thus, a naïve grounding which replaces the variables with all the constants appearing in the program (thus producing the full instantiation) is undesirable from a computational point of view. Indeed, most of the ground atoms appearing in the full instantiation are not derivable from the program rules, and all generated ground rules containing these atoms in the positive bodies are useless for answer set computation. For instance, consider the following program:

```

c(1, 2).
a(X) | b(Y) :- c(X, Y).
  
```

The full instantiation of the only rule appearing in the program contains four ground instances:

```

a(1) | b(1) :- c(1, 1).
a(2) | b(1) :- c(2, 1).
a(2) | b(2) :- c(2, 2).
a(1) | b(2) :- c(1, 2).
  
```

However, the first three ground rules are useless. They will never be applicable because their bodies contain atoms $c(1, 1)$ and $c(2, 1)$, and $c(2, 2)$ that are not derivable from the program (they do not appear in the head of any rule).

ASP grounders, like GrinGo or the DLV instantiator, employ smart procedures that are geared toward efficiently producing a ground program that is considerably smaller than the full instantiation but preserves the semantics. In the following, we first give an informal description of the grounding computa-

tion. Then we introduce the problem of dealing with function symbols, which may lead to infinite groundings. Finally we overview some optimization strategies.

The Instantiation Procedure

In this subsection, we provide a description of the basic instantiation procedure, which is adopted by the most popular grounders, GrinGo and the DLV instantiator. For clarity, the description is informal, and presents a simplified version of the actual instantiation strategy. For instance, we do not take into account extensions of the basic language like choice rules or aggregates (Lifschitz 2016; Alviano and Leone 2015, 2016). Full details can be found in the work by Faber, Leone, and Perri (2012) and Gebser et al. (2011).

The core of the grounding phase is the process of rule instantiation. Given a rule r and a set of ground atoms S , which represents the extensions of the predicates, it generates the ground instances of r . Such a task can be performed by iterating on the body literals looking for possible substitutions for their variables. Grounders impose a safety condition, which requires that each rule variable appear also in a positive body literal. Thus, for the instantiator, it is enough to have a substitution for the variables occurring in the positive literals.

To clarify this process, consider the following (nonground) rule:

$$a(X) \mid b(Y) \text{ :- } p(X, Z), q(Z, Y).$$

Now, assume that the set of extensions $S = \{p(1, 2), q(2, 1), q(2, 3)\}$ is given. Then, the instantiation starts by looking for a ground atom in S matching with $p(X, Z)$. Therefore $p(X, Z)$ is matched with $p(1, 2)$ and the substitution for X and Z is propagated to the other body literals, thus leading to the partially ground rule body $p(1, 2), q(2, Y)$. Then, $q(2, Y)$ is instantiated with the matching ground atom $q(2, 1)$ and a ground rule $a(1) \mid b(1) \text{ :- } p(1, 2), q(2, 1)$ is generated. Now, in order to find other possible rule instances, a backtracking step is performed, the binding for variable Y is restored and a new match for $q(2, Y)$ is searched, finding $q(2, 3)$. The new match is applied, leading to another rule instance $a(1) \mid b(3) \text{ :- } p(1, 2), q(2, 3)$. Then, the process goes on, by backtracking again to $q(2, Y)$, and then to $p(X, Z)$, because there are no more matches for $q(2, Y)$. Given that also no further matches are possible for $p(X, Z)$, the instantiation of the rule terminates, producing only two ground rules:

$$a(1) \mid b(1) \text{ :- } p(1, 2), q(2, 1).$$

$$a(1) \mid b(3) \text{ :- } p(1, 2), q(2, 3).$$

Roughly, the body literals are instantiated from left to right, starting from the first one. The instantiation of the generic body literal L consists in searching in S for a ground atom A matching with L ; if such a matching is found, then the variables in L are bound with the constants in A , the substitution is propagated to the other body literals, and the next literal in

the body is considered. If such a matching atom is not found, a backtracking step to a previous literal L' is performed, some variable bindings are restored, and the process goes on by looking for another matching for L' . When all body literals have been instantiated, an instance for the rule r is found and the process continues by backtracking again to some previous literal, in order to find other substitutions. A crucial aspect of this process is how the set of ground atoms S containing the extensions of the predicates is computed. When a program is given as input to a grounder, it usually contains also a set of ground atoms, called *Facts*. It constitutes the starting point of the computation. In other words, initially $S = \text{Facts}$. During instantiation, the set S is expanded with the ground atoms occurring in the head of the newly generated ground rules. For instance, in the previous example, the ground atoms $a(1)$ and $b(1)$ are added to S and they will possibly be used for the instantiation of other rules. Thus, the extensions of the predicates are built dynamically. In order to guarantee the generation of all useful ground instances a particular evaluation order should be followed. If a rule r_1 defines (that is, has in the head) a predicate p , and another rule r_2 contains p in the positive body, then r_1 has to be evaluated before r_2 since r_1 produces ground atoms needed for instantiating r_2 . Complying with such evaluation orders ensures that the produced ground program has the same answer sets of the full instantiation, but is possibly smaller (Faber, Leone, and Perri 2012).

To produce proper evaluation orders, grounders make use of structural information provided by a directed graph, called *Dependency Graph*, that describes how predicates depend on each other. This graph induces a partition of the input program into subprograms, associated with the strongly connected components, and a topological ordering over them. The subprograms are instantiated one at a time starting from the ones associated with the lowest components in the topological ordering.

Recursive rules within a subprogram, that is, rules where some body predicate depends, directly or transitively, on a predicate in the head, are instantiated according to a seminaive database technique (Ullman 1988). Their evaluation produces ground atoms needed for their own evaluation, thus, several iterations are performed, until a fixpoint is reached. At each iteration, for the predicates involved in the recursion, only the ground atoms newly derived during the previous iteration are taken into account.

To illustrate this, consider the following problem, called *Reachability*: Given a finite directed graph, compute all pairs of nodes (a, b) such that b is reachable from a through a nonempty sequence of arcs. This problem can be encoded by the following ASP program:

$$\text{reach}(X, Y) \text{ :- } \text{arc}(X, Y).$$

$$\text{reach}(X, Y) \text{ :- } \text{arc}(X, U), \text{reach}(U, Y).$$

The set of arcs is represented by the binary relation *arc*. A fact *arc(a, b)* means that the graph contains an arc from *a* to *b*; the set of nodes is not explicitly represented.

The program computes a binary relation *reach* containing all facts *reach(a, b)* such that *b* is reachable from *a* through the arcs of the input graph *G*. In particular, the first (nonrecursive) rule states that *b* is directly reachable from *a*, if there is an arc from *a* to *b*; while the second (recursive) rule states that *b* is transitively reachable from *a*, if there is a path in the graph from *a* to *b*.

The instantiation of this program is performed by first evaluating the nonrecursive rule on the set *S* containing the arcs. Assuming that $S = \{ \text{arc}(1, 2), \text{arc}(2, 3), \text{arc}(3, 4) \}$ three ground instances are produced:

```
reach(1, 2) :- arc(1, 2).
reach(2, 3) :- arc(2, 3).
reach(3, 4) :- arc(3, 4).
```

The ground atoms *reach(1, 2)*, *reach(2, 3)*, and *reach(3, 4)* are added to set *S* and the evaluation of the recursive rule starts. The first iteration is performed, producing rules

```
reach(1, 3) :- arc(1, 2), reach(2, 3).
reach(2, 4) :- arc(2, 3), reach(3, 4).
```

Then, *reach(1, 3)*, *reach(2, 4)* are added to *S* and another iteration starts.

To avoid duplicate rules, for the recursive predicate *reach*, only the two newly generated ground atoms are used, producing:

```
reach(1, 4) :- arc(1, 2), reach(2, 4).
```

Now, *reach(1, 4)* is added to *S*. Another iteration is performed. Nothing new can be produced. The fix-point is reached and the evaluation terminates.

Optimizations

Substantial effort has been spent on sophisticated algorithms and optimization techniques aimed at improving the performance of the instantiation process. In the following we briefly recall the most relevant ones.

The dynamic magic sets technique (Alviano et al. 2012) is a rewriting-based optimization strategy used by the DLV system. It extends the Magic Sets technique originally defined for standard Datalog for optimizing query answering over logic programs. Given a query, the Magic Sets technique rewrites the input program to identify a subset of the program instantiation which is sufficient for answering the query. The restriction of the instantiation is obtained by means of additional “magic” predicates, whose extensions represent relevant atoms with regard to the query. Dynamic Magic Sets, specifically conceived for disjunctive programs, inherit the benefits provided by standard magic sets and additionally allow to exploit the information provided by the magic predicates also during the answer set search.

Magic sets turned out to be very useful in many application domains, even on some co-NP complete problems like consistent query answering (Manna, Ricca, and Terracina 2015).

Other techniques have been developed for optimizing the rule instantiation task (Faber, Leone, and Perri 2012). In particular, since rule instantiation is essentially performed by evaluating the relational join of the positive body literals, an optimal ordering of literals in the body is a key issue for the efficiency of the procedure, just like for join computation. Thus, an efficient body reordering criterion specifically conceived for the rule instantiation task has been proposed. Moreover, a backjumping algorithm has been developed (Perri et al. 2007), which reduces the size of the ground programs, avoiding the generation of useless rules, but fully preserving the semantics.

In the last few years, in order to make use of modern multicore and multiprocessor computers, a parallel instantiator has been developed. It is based on a number of strategies allowing for the concurrent evaluation of parts of the program, and is endowed with advanced mechanisms for dealing with load balancing and granularity control (Perri, Ricca, and Sirianni 2013).

Dealing with Function Symbols

Function symbols are widely recognized as an important feature for ASP. They increase the expressive power and in some cases improve the modeling capabilities of ASP, allowing the support of complex terms like lists, and set terms. Functions can also be employed to represent, through skolemization, existential quantifiers, which are receiving an increasing attention in the logic programming and database communities (Gottlob, Manna, and Pieris 2015). However, the presence of function symbols within ASP programs has a strong impact on the grounding process, which might even not terminate. Consider, for instance, the program:

```
p(0).
p(f(X)) :- p(X).
```

The instantiation is infinite; indeed the grounding of the recursive rule, at the first iteration adds to the set of extensions *S* the ground atom *p(f(0))*, which is used in the next iteration, producing *p(f(f(0)))* and so on. Despite this, grounders like the one in DLV and GrinGo are a way to deal with recursive function symbols and guarantee termination whenever the program belongs to the class of the so called finitely ground programs (Calimeri et al. 2008). Intuitively, for each program *P* in this class, there exists a finite ground program *P'* having exactly the same answer sets as *P*. Program *P'* is computable for finitely ground programs, thus answer sets of *P* are computable as well. Notably, each computable function can be expressed by a finitely ground program; membership in this class is not decidable, but it has been proven to be semidecidable (Calimeri et al. 2008).

For applications in which termination needs to be guaranteed a priori, the ASP grounders can make use of a preprocessor implementing a decidable check, which allows the user to statically recognize whether the input program belongs to a smaller subclass of the finitely ground programs (Syrjänen 2001; Gebser, Schaub, and Thiele 2007; Lierler and Lifschitz 2009; Calimeri et al. 2008). For instance, the grounder of DLV is endowed with a checker (which can also be disabled) for recognizing argument-restricted programs (Lierler and Lifschitz 2009). Earlier versions of GrinGo, in order to guarantee finiteness, accepted input programs with a domain restriction, namely λ -restricted programs (Gebser, Schaub, and Thiele 2007). From series 3, GrinGo removed domain restrictions and the responsibility to check whether the input program has a finite grounding is left to the user.

Solving

Modern ASP solvers rely upon advanced conflict-driven search procedures, pioneered in the area of satisfiability testing (SAT; [Biere et al. 2009]).³ Conflicts are analyzed and recorded, decisions are taken in view of conflict scores, and back-jumps are directed to the origin of a conflict.

While the general outline of search in ASP is arguably the same as in SAT, the extent of ASP requires a much more elaborate approach. First, the stable model's semantics enforces that atoms are not merely true but provably true (Lifschitz 2016). Second, the rich modeling language of ASP comes with complex language constructs. In particular, disjunction in rule heads and nonmonotone aggregates lead to an elevated level of computational complexity, which imposes additional search efforts. Finally, ASP deals with various reasoning modes. Apart from satisfiability testing, this includes enumeration, projection, intersection, union, and (multiobjective) optimization of answer sets, and moreover combinations of them, for instance, the intersection of all optimal models.⁴ The first two issues bring about additional inferences, the latter require flexible solver architectures.

The restriction of modern SAT solvers to propositional formulas in conjunction normal form allows for reducing inferences to unit propagation along with the usual choice operations. In contrast, traditional ASP solving deals with an abundance of different inferences for propagation, which makes a direct adaption of conflict-driven search procedures virtually impossible. The key idea is thus to map inferences in ASP onto unit propagation on nogoods⁵ (Gebser, Kaufmann, and Schaub 2012a), which traces back to a characterization of answer sets in propositional logic (Lin and Zhao 2004). Let us illustrate this by program P (thereby restricting ourselves to normal rules):⁶

$$P = \left\{ \begin{array}{l} a : \text{-not}b, \quad b : \text{-not}a, \\ x : \text{-}a, \text{not}c, \quad x : \text{-}y, \\ y : \text{-}x, b \end{array} \right\}$$

Interpreting this program in propositional logic results in the set $RF(P)$ of implications:

$$RF(P) = \left\{ \begin{array}{l} a \leftarrow \text{-}b, \quad b \leftarrow \text{-}a, \\ x \leftarrow a \wedge \text{-}c \vee y, \\ y \leftarrow x \wedge b \end{array} \right\}$$

Note that we replaced default negation not by classical negation \neg and combined both rules with head x while leaving the direction of the implications untouched (for readability). Now, the set $RF(P)$ has twelve classical models, many of which contain atoms not supported by any rule. (This is important because the stable models semantics insists on provably true atoms.) For instance, c is not supported by any rule, as is b whenever a is true as well.

Models containing unsupported atoms are eliminated by turning the implications in $RF(P)$ into equivalences (Clark 1978). Doing so for each atom yields the set $CF(P)$ of equivalences:

$$CF(P) = \left\{ \begin{array}{l} a \leftrightarrow \text{-}b, \quad b \leftrightarrow \text{-}a, \\ x \leftrightarrow a \wedge \text{-}c \vee y, \\ y \leftrightarrow x \wedge b, \quad c \leftrightarrow \perp \end{array} \right\}$$

This strengthening results in three models of $CF(P)$, one entailing atom b only, another making b, x, y true, and finally one in which a, x hold, respectively. The first two models differ in making both x and y true or not. A closer look at the original program P reveals that x and y support each other in a circular way. Whether or not such a circular derivation is harmful depends upon the existence of a valid external support (Lin and Zhao 2004), provided by an applicable rule whose head is in the loop but none of its positive antecedents belongs to it. In our case, this can be accomplished by the formula in $LF(P)$:

$$LF(P) = \left\{ (x \vee y) \rightarrow a \wedge \text{-}c \right\}$$

The formula expresses that an atom in the loop (consisting of x and y) can only be true if an external support of x or y is true. Here the only external support is provided by rule $x \leftarrow a, \text{not}c$ in P , as reflected by the consequent in $LF(P)$. That is, x or y can only be true if the latter rule applies. Since no other loops occur in P , the set $CF(P) \cup LF(P)$ provides a characterization of P 's answer sets (Lin and Zhao 2004), one making atom b true and another a, x .

Note that in general the size of $CF(P)$ is linear in that of a program P , whereas the size of $LF(P)$ may be exponential in P (Lifschitz and Razborov 2006). Fortunately, satisfaction of $LF(P)$ can be tested in linear time for logic programs facing no elevated complexity (previously discussed), otherwise this test is co-

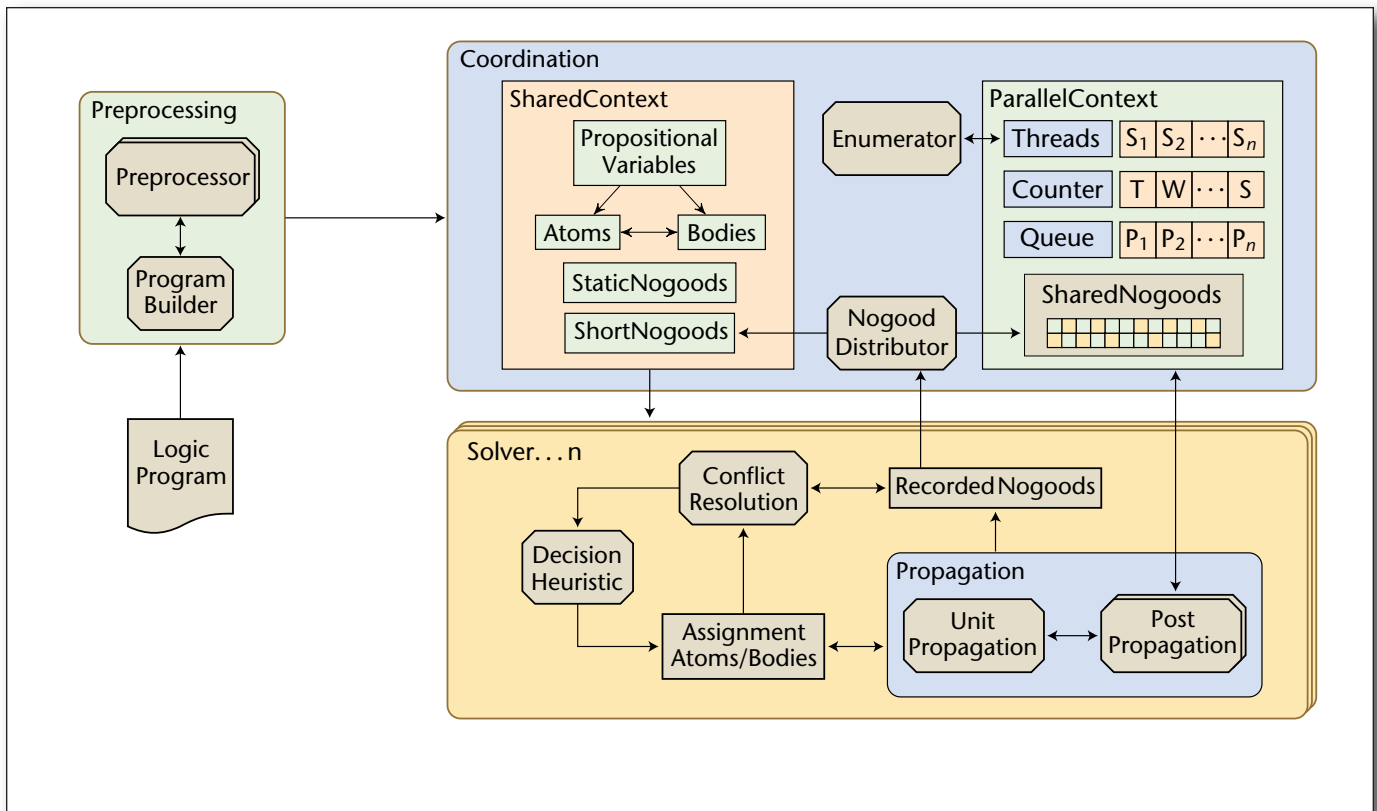


Figure 2. The Multithreaded Architecture of the ASP Solver Clasp.

NP-complete (Leone, Rullo, and Scarcello 1997).

The translation of programs into nogoods employed by modern ASP solvers follows the above characterization but takes the space issue into account. Given a program P , the nogoods expressing $CF(P)$ are explicitly represented in an ASP solver, while the ones in $LF(P)$ are only made explicit upon violation. This violation is detected by so-called unfounded set algorithms (Leone, Rullo, and Scarcello 1997; Gebser et al. 2012). Although we do not detail this here, we mention that aggregates are treated in a similar way by dedicated mechanisms unless they can be translated into nogoods in a feasible way (Gebser et al. 2009). Finally, let us make this more concrete by looking at the system architecture of clasp, depicted in figure 2.

The preprocessing component takes a (disjunctive) logic program and translates it into an internal representation. This is done in several steps. First, the given program, P , is simplified by semantic preserving translations as well as equivalence detection (Gebser et al. 2008). The simplified program P' is then translated into nogoods expressing $CF(P')$, which are subject to clausal simplifications adapted from corresponding SAT techniques. The resulting static nogoods are kept in the shared context component, as are parts of the dependency graph of P in

order to reconstruct members of $LF(P)$ on demand. Often more than three quarters of the nogoods obtained from $CF(P')$ are binary or ternary. Hence, such short nogoods are stored in dedicated data structures (and shared during parallel solving). Each solver instance implements a conflict-driven search procedure, as sketched at the outset of this section. Of particular interest is propagation, distinguishing between unit and post propagation. The former computes a fixed point of unit propagation. More elaborate propagation mechanisms can be added through post propagators. For instance, for programs with loops, this list contains a post propagator implementing the unfounded set checking procedure. Similarly, clasp's extension with constraint processing, clingcon (Ostrowski and Schaub 2012), as well as dlhex (Eiter et al. 2006) use its post propagation mechanism to realize additional theory-specific propagations. The parallel execution of clasp allows for search space splitting as well as running competitive strategies. In both cases, learned conflict nogoods (as well as bounds in case of optimization) are exchanged between solver instances, each of which can be configured individually (see Gebser, Kaufmann, and Schaub [2012b] for details on multithreading). Finally, the enumerator is in charge of handling the various reasoning modes; once a solver

finds a model, the enumerator tells it how to continue. This can be done by finding a next model in case of enumeration, or a better model in case of optimization.

Conclusion

Answer set programming combines a high-level modeling language with effective grounding and solving technology. This materializes in off-the-shelf ASP systems, whose grounding and solving engines can be used as black-box systems with standardized interfaces. Also, ASP is highly versatile by offering various complex language constructs and reasoning modes. As a side effect, many ASP solvers can also be used for MAX-SAT⁷, SAT, and PB⁸ solving. As a consequence, ASP faces a growing range of applications, as detailed by Erdem, Gelfond, and Leone (2016).

Acknowledgments

The first and last author were partially funded by DFG grants SCHA 550/8 and SCHA 550/9. The second and third author were partially supported by MIUR under PON project SI-LAB BA2KNOW Business Analytics to Know, and by Regione Calabria, programme POR Calabria FESR 2007-2013, projects ITravel PLUS and KnowRex: Un sistema per il riconoscimento e lestrazione di conoscenza.

Notes

1. Satisfiability testing.
2. Mixed integer programming.
3. This technology is usually referred to as conflict-driven clause learning.
4. Actually, this is a frequent reasoning mode used in under-specified application domains such as bioinformatics (Erdem, Gelfond, and Leone 2016).
5. Nogoods express inadmissible assignments (Dechter 2003).
6. $RF(P)$, $CF(P)$, and $LF(P)$ stand for the rule, completion, and loop formulas of P .
7. Maximum satisfiability problem.
8. Pseudo-Boolean.

References

- Alviano, M., and Leone, N. 2015. Complexity and Compilation of GZ-Aggregates in Answer Set Programming. *Theory and Practice of Logic Programming* 15(4-5): 574–587. dx.doi.org/10.1017/S147106841500023X
- Alviano, M., and Leone, N. 2016. On the Properties of GZ-Aggregates in Answer Set Programming. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16)*. Palo Alto, CA: AAAI Press.
- Alviano, M.; Dodaro, C.; Leone, N.; and Ricca, F. 2015. Advances in WASP. In *Proceedings of the Thirteenth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'15)*, ed. F. Calimeri, G. Ianni, and M. Truszczyński, 40–54. Berlin: Springer. dx.doi.org/10.1007/978-3-319-23264-5_5
- Alviano, M.; Faber, W.; Greco, G.; and Leone, N. 2012. Magic Sets for Disjunctive Datalog Programs. *Artificial Intelligence* 187: 156–192. dx.doi.org/10.1016/j.artint.2012.04.008
- Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds. 2009. *Handbook of Satisfiability*. Amsterdam: IOS Press.
- Calimeri, F.; Cozza, S.; Ianni, G.; and Leone, N. 2008. Computable Functions in ASP: Theory and Implementation. In *Proceedings of the Twenty-Fourth International Conference on Logic Programming (ICLP'08)*, ed. M. Garcia de la Banda and E. Pontelli, 407–424. Berlin: Springer. dx.doi.org/10.1007/978-3-540-89982-2_37
- Clark, K. 1978. Negation as Failure. In *Logic and Data Bases*, ed. H. Gallaire and J. Minker, 293–322. New York: Plenum Press. dx.doi.org/10.1007/978-1-4684-3384-5_11
- Dantsin, E.; Eiter, T.; Gottlob, G.; and Voronkov, A. 2001. Complexity and Expressive Power of Logic Programming. *ACM Computing Surveys* 33(3): 374–425. dx.doi.org/10.1145/502807.502810
- Dechter, R. 2003. *Constraint Processing*. San Francisco: Morgan Kaufmann Publishers.
- Eiter, T.; Ianni, G.; Schindlauer, R.; and Tompits, H. 2006. DLVHEX: A Prover for Semantic-Web Reasoning under the Answer-Set Semantics. In *Proceedings of the International Conference on Web Intelligence (WI'06)*, 1073–1074. Los Alamitos, CA: IEEE Computer Society. dx.doi.org/10.1109/wi.2006.64
- Erdem, E.; Gelfond, M.; and Leone, N. 2016. Applications of ASP. *AI Magazine*. 37(3).
- Faber, W.; Leone, N.; and Perri, S. 2012. The Intelligent Grounder of DLV. In *Correct Reasoning: Essays on Logic-Based AI in Honour of Vladimir Lifschitz*, ed. E. Erdem, J. Lee, Y. Lierler, and D. Pearce. Berlin: Springer. 247–264. dx.doi.org/10.1007/978-3-642-30743-0_17
- Gebser, M., and Schaub, T. 2016. Modeling and Language Extensions. *AI Magazine* 37(3).
- Gebser, M.; Kaminski, R.; Kaufmann, B.; and Schaub, T. 2009. On the Implementation of Weight Constraint Rules in Conflict-Driven ASP Solvers. In *Proceedings of the Twenty-Fifth International Conference on Logic Programming (ICLP'09)*, ed. P. Hill and D. Warren, 250–264. Berlin: Springer. dx.doi.org/10.1007/978-3-642-02846-5_23
- Gebser, M.; Kaminski, R.; Kaufmann, B.; and Schaub, T. 2012. *Answer Set Solving in Practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. San Rafael, CA: Morgan and Claypool Publishers.
- Gebser, M.; Kaminski, R.; König, A.; and Schaub, T. 2011. Advances in Gringo Series 3. In *Proceedings of the Eleventh International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'11)*, ed. J. Delgrande, and W. Faber, 345–351. Berlin: Springer. dx.doi.org/10.1007/978-3-642-20895-9_39
- Gebser, M.; Kaufmann, B.; and Schaub, T. 2012a. Conflict-Driven Answer Set Solving: From Theory to Practice. *Artificial Intelligence* 187–188: 52–89. dx.doi.org/10.1016/j.artint.2012.04.001
- Gebser, M.; Kaufmann, B.; and Schaub, T. 2012b. Multi-Threaded ASP Solving with Clasp. *Theory and Practice of Logic Programming* 12(4–5): 525–545. dx.doi.org/10.1017/S1471068412000166
- Gebser, M.; Kaufmann, B.; Neumann, A.; and Schaub, T. 2008. Advanced Preprocessing for Answer Set Solving. In *Proceedings of the Eighteenth European Conference on Artificial Intelligence (ECAI'08)*, ed. M. Ghallab, C. Spyropoulos, N. Fakotakis, and N. Avouris, 15–19. Amsterdam: IOS Press.

- Gebser, M.; Schaub, T.; and Thiele, S. 2007. Gringo: A New Grounder for Answer Set Programming. In *Logic Programming and Nonmonotonic Reasoning, 9th International Conference*, Lecture Notes in Computer Science 4483, ed. C. Baral, G. Brewka, and J. Schlipf, 266–271. Berlin: Springer. dx.doi.org/10.1007/978-3-540-72200-7_24
- Giunchiglia, E.; Lierler, Y.; and Maratea, M. 2006. Answer Set Programming Based on Propositional Satisfiability. *Journal of Automated Reasoning* 36(4): 345–377. dx.doi.org/10.1007/s10817-006-9033-2
- Gottlob, G.; Manna, M.; and Pieris, A. 2015. Polynomial Rewritings for Linear Existential Rules. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-15)*, 2992–2998. Palo Alto, CA: AAAI Press.
- Janhunen, T.; Niemelä, I.; and Sevalnev, M. 2009. Computing Stable Models via Reductions to Difference Logic. In *Logic Programming and Nonmonotonic Reasoning, 9th International Conference, LPNMR 2007*, Lecture Notes in Computer Science 4483, ed. E. Erdem, F. Lin, and T. Schaub, 142–154. Berlin: Springer. dx.doi.org/10.1007/978-3-642-04238-6_14
- Leone, N.; Pfeifer, G.; Faber, W.; Eiter, T.; Gottlob, G.; Perri, S.; and Scarcello, F. 2006. The DLV System for Knowledge Representation and Reasoning. *ACM Transactions on Computational Logic* 7(3): 499–562. dx.doi.org/10.1145/1149114.1149117
- Leone, N.; Rullo, P.; and Scarcello, F. 1997. Disjunctive Stable Models: Unfounded Sets, Fixpoint Semantics, and Computation. *Information and Computation* 135(2): 69–112. dx.doi.org/10.1006/inco.1997.2630
- Lierler, Y., and Lifschitz, V. 2009. One More Decidable Class of Finitely Ground Programs. In *Logic Programming, 25th International Conference*, Lecture Notes in Computer Science 5649, 489–493. Berlin: Springer. dx.doi.org/10.1007/978-3-642-02846-5_40
- Lifschitz, V. 2016. Answer Sets and the Language of Answer Set Programming. *AI Magazine* 37(3).
- Lifschitz, V., and Razborov, A. 2006. Why Are There So Many Loop Formulas? *ACM Transactions on Computational Logic* 7(2): 261–268. dx.doi.org/10.1145/1131313.1131316
- Lin, F., and Zhao, Y. 2004. ASSAT: Computing Answer Sets of a Logic Program by SAT Solvers. *Artificial Intelligence* 157(1–2): 115–137. dx.doi.org/10.1016/j.artint.2004.04.004
- Liu, G.; Janhunen, T.; and Niemelä, I. 2012. Answer Set Programming via Mixed Integer Programming. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2012)*, ed. G. Brewka, T. Eiter, and S. McIlraith, 32–42. Palo Alto, CA: AAAI Press.
- Manna, M.; Ricca, F.; and Terracina, G. 2015. Taming Primary Key Violations to Query Large Inconsistent Data via ASP. *Theory and Practice of Logic Programming* 15(4–5): 696–710. dx.doi.org/10.1017/S1471068415000320
- Maratea, M.; Pulina, L.; and Ricca, F. 2014. A Multi-Engine Approach to Answer-Set Programming. *Theory and Practice of Logic Programming* 14(6): 841–868. dx.doi.org/10.1017/S1471068413000094
- Ostrowski, M., and Schaub, T. 2012. ASP Modulo CSP: The Clingcon System. *Theory and Practice of Logic Programming* 12(4–5): 485–503. dx.doi.org/10.1017/S1471068412000142
- Perri, S.; Ricca, F.; and Sirianni, M. 2013. Parallel Instantiation of ASP Programs: Techniques and Experiments. *Theory and Practice of Logic Programming* 13(2): 253–278. dx.doi.org/10.1017/S1471068411000652
- Perri, S.; Scarcello, F.; Catalano, G.; and Leone, N. 2007. Enhancing DLV Instantiator by Backjumping Techniques. *Annals of Mathematics and Artificial Intelligence* 51(2–4): 195–228. dx.doi.org/10.1007/s10472-008-9090-9
- Simons, P.; Niemelä, I.; and Sooinen, T. 2002. Extending and Implementing the Stable Model Semantics. *Artificial Intelligence* 138(1–2): 181–234. dx.doi.org/10.1016/S0004-3702(02)00187-X
- Syrjänen, T. 2001. Omega-Restricted Logic Programs. In *Logic Programming and Nonmonotonic Reasoning, 6th International Conference*, Lecture Notes in Computer Science 2173, 267–279, ed. T. Eiter, F. Faber, and M. Truszczynski. Berlin: Springer. dx.doi.org/10.1007/3-540-45402-0_20
- Ullman, J. 1988. *Principles of Database and Knowledge-Base Systems*. Rockville, MD: Computer Science Press.
- Wittocx, J.; Mariën, M.; and Denecker, M. 2010. Grounding FO and FO(ID) with Bounds. *Journal of Artificial Intelligence Research* 38: 223–269.

Benjamin Kaufmann is a postdoctoral researcher at the University of Potsdam, Germany. He mainly works on efficient search-based reasoning and optimization techniques for Boolean satisfiability and related approaches. In particular, he is the lead developer of clasp, a state-of-the-art solving system in answer set programming.

Nicola Leone is a professor of computer science at University of Calabria, where he heads the Department of Mathematics and Computer Science and leads the AI Lab. He was professor of database systems at Vienna University of Technology until 2000. He is internationally renowned for his research on knowledge representation, answer set programming (ASP), and database theory, and for the development of DLV, a state-of-the-art ASP system that is popular worldwide. He has published more than 250 papers in prestigious conferences and journals, and has more than 8000 citations, with h-index 46. He is a fellow of ECCAI (now EurAI) and recipient of a Test of Time award from the Association for Computing Machinery.

Simona Perri is an associate professor at University of Calabria, Italy. She received her MSc in mathematics (2000) and a Ph.D. in computer science and mathematics (2005) from the University of Calabria, Italy. She is a member of the team that designed and maintains DLV, one of the major answer set programming systems, and is cofounder of DLVSystem Ltd, a spin-off company of the University of Calabria.

Torsten Schaub is university professor at the University of Potsdam, Germany, and holds an international chair at Inria Rennes, France. He is a fellow of ECCAI and the current president of the Association of Logic Programming. His current research focus, on answer set programming (ASP) and its applications, materializes at potassco.sourceforge.net, the home of the open source project Potassco bundling software for ASP developed at Potsdam.