

# *Multi-shot ASP solving with clingo\**

MARTIN GEBSER, ROLAND KAMINSKI  
and BENJAMIN KAUFMANN

*University of Potsdam, Germany*

TORSTEN SCHAUB

*INRIA Rennes, France  
University of Potsdam, Germany*

*submitted 2 April 2018; accepted 3 April 2018*

---

## Abstract

We introduce a new flexible paradigm of grounding and solving in Answer Set Programming (ASP), which we refer to as multi-shot ASP solving, and present its implementation in the ASP system *clingo*. Multi-shot ASP solving features grounding and solving processes that deal with continuously changing logic programs. In doing so, they remain operative and accommodate changes in a seamless way. For instance, such processes allow for advanced forms of search, as in optimization or theory solving, or interaction with an environment, as in robotics or query answering. Common to them is that the problem specification evolves during the reasoning process, either because data or constraints are added, deleted, or replaced. This evolutionary aspect adds another dimension to ASP since it brings about state changing operations. We address this issue by providing an operational semantics that characterizes grounding and solving processes in multi-shot ASP solving. This characterization provides a semantic account of grounder and solver states along with the operations manipulating them. The operative nature of multi-shot solving avoids redundancies in relaunching grounder and solver programs and benefits from the solver’s learning capacities. *clingo* accomplishes this by complementing ASP’s declarative input language with control capacities. On the declarative side, a new directive allows for structuring logic programs into named and parameterizable subprograms. The grounding and integration of these subprograms into the solving process is completely modular and fully controllable from the procedural side. To this end, *clingo* offers a new application programming interface that is conveniently accessible via scripting languages. By strictly separating logic and control, *clingo* also abolishes the need for dedicated systems for incremental and reactive reasoning, like *iclingo* and *oclingo*, respectively, and its flexibility goes well beyond the advanced yet still rigid solving processes of the latter.

---

## 1 Introduction

Standard Answer Set Programming (ASP; Baral (2003)) follows a one-shot process in computing stable models of logic programs. This view is best reflected

\* We are grateful to Evgenii Balai, Javier Romero, and Adam Smith for many fruitful discussions on the paper. This work was partially funded by DFG Grant SCHA 550/9.

by the input/output behavior of monolithic ASP systems like *dlv* (Leone et al. 2006) and (original) *clingo* (Gebser et al. 2011) that take a logic program and output its stable models. Internally, however, both follow a fixed two-step process. First, a grounder generates a (finite) propositional representation of the input program. Then, a solver computes the stable models of the propositional program. This rigid process stays unchanged when grounding and solving with separate systems. In fact, up to series 3, *clingo* was a mere combination of the grounder *gringo* and the solver *clasp*. Although more elaborate reasoning processes are performed by the extended systems *iclingo* (Gebser et al. 2008) and *oclingo* (Gebser et al. 2011) for incremental and reactive reasoning, respectively, they also follow a pre-defined control loop evading any user control. Beyond this, however, there is substantial need for specifying flexible reasoning processes, for instance, when it comes to interactions with an environment (as in assisted living, robotics, or with users), advanced search (as in multi-objective optimization, planning, theory solving, or heuristic search), or recurrent query answering (as in hardware analysis and testing or stream processing). Common to all these advanced forms of reasoning is that the problem specification evolves during the respective reasoning processes, either because data or constraints are added, deleted, or replaced.

For mastering such complex reasoning processes, we propose the paradigm of *multi-shot ASP solving* in order to deal with continuously changing logic programs. In contrast to the traditional single-shot approach, where an ASP system takes a logic program, computes its answer sets, and exits, the idea is to consider evolving grounding and solving processes. Such processes lead to operative ASP systems that possess an internal state that can be manipulated by certain operations. Such operations allow for adding, grounding, and solving logic programs as well as setting truth values of (external) atoms. The latter does not only provide a simple means for incorporating external input but also for enabling or disabling parts of the current logic program. These functionalities allow for dealing with changing logic programs in a seamless way. To capture multi-shot solving processes, we introduce an operational semantics centered upon a formal definition of an ASP system state along with its (state changing) operations. Such a state reflects the relevant information gathered in the system's grounder and solver components. This includes (i) a collection of non-ground logic programs subject to grounding, (ii) the ground logic programs currently held by the solver, (iii) and a truth assignment of externally defined atoms. Changing such a state brings about several challenges absent in the single-shot case, among them, contextual grounding and logic program composition.

Given that the theoretical foundations of multi-shot solving are a means to an end, we interleave their presentation with the corresponding features of ASP system *clingo*. This new generation of *clingo*<sup>1</sup> offers novel high-level constructs for

<sup>1</sup> Multi-shot solving was introduced with the *clingo* 4 series by Gebser et al. (2014). However, we describe its functionalities in the context of the current *clingo* 5 series. The advance from series 4 to 5 only smoothed some multi-shot related interfaces but left the principal functionality unaffected.

realizing multi-shot ASP solving. This is achieved within a single ASP grounding and solving process that avoids redundancies otherwise caused by relaunching grounder and solver programs and benefits from the learning capacities of modern ASP solvers. To this end, *clingo* complements ASP's declarative input language by manifold control capacities. The latter are provided by an imperative application programming interface (API) implemented in C. Corresponding bindings for Python and Lua are available and can also be embedded into *clingo*'s input language (via the `#script` directive). On the declarative side, *clingo* offers a new directive `#program` that allows for structuring logic programs into named and parametrizable subprograms. The grounding and integration of these subprograms into the solving process is completely modular and fully controllable from the procedural side. For exercising control, the latter benefits from a dedicated library furnished by *clingo*'s API, which does not only expose grounding and solving functions but moreover allows for continuously assembling the solver's program. This can be done in combination with externally controllable atoms that allow for enabling or disabling rules. Such atoms are declared by the `#external` directive. Hence, by strictly separating logic and control, *clingo* abolishes the need for special-purpose systems for incremental and reactive reasoning, like *iclingo* and *oclingo*, respectively, and its flexibility goes well beyond the advanced yet still rigid grounding and solving processes of such systems. In fact, *clingo*'s multi-shot solving capabilities rather enable users to engineer novel forms of reasoning, as we demonstrate by four case studies.

The rest of the paper is organized as follows. Section 2 provides a brief account of formal preliminaries. Section 3 gives an informal overview on the new features of *clingo* in order to pave the way for their formal underpinnings presented in Section 4. There, we lay the formal foundations of multi-shot solving and present its aforementioned operational semantics. In Section 5, we illustrate the power of multi-shot ASP solving in several use cases and highlight some features of interest. We use Python throughout the paper to illustrate the multi-shot functionalities of *clingo*'s API. Further API-related aspects are described in Section 6. Section 7 gives an empirical analysis of some selected features of multi-shot solving with *clingo*. Finally, we relate our approach to the literature in Section 8, before we conclude in Section 9.

## 2 Formal preliminaries

A (normal<sup>2</sup>) rule  $r$  is an expression of the form

$$a_0 \leftarrow a_1, \dots, a_m, \sim a_{m+1}, \dots, \sim a_n,$$

where  $a_i$ , for  $0 \leq m \leq n$ , is an *atom* of the form  $p(t_1, \dots, t_k)$ ,  $p$  is a predicate symbol of arity  $k$ , also written as  $p/k$ , and  $t_1, \dots, t_k$  are terms, built from constants, variables, and functions. Letting  $h(r) = a_0$ ,  $B(r) = \{a_1, \dots, a_m, \sim a_{m+1}, \dots, \sim a_n\}$ ,  $B(r)^+ = \{a_1, \dots, a_m\}$ ,

<sup>2</sup> For the sake of simplicity, we confine our formal elaboration to normal logic programs. Multi-shot solving with *clingo* also works with disjunctive logic programs.

and  $B(r)^- = \{a_{m+1}, \dots, a_n\}$ , we also denote  $r$  by  $h(r) \leftarrow B(r)$ . A rule is called fact, whenever  $B(r) = \emptyset$ . A (normal) *logic program*  $P$  is a set (or list) of rules (depending on whether the order of rules matters or not). We write  $H(P) = \{h(r) \mid r \in P\}$  and  $A(P) = H(P) \cup \bigcup_{r \in P} (B(r)^+ \cup B(r)^-)$  to denote the set of all head atoms and atoms occurring in  $P$ , respectively. A term, atom, rule, or program is *ground* if it does not contain variables.

We denote the set of all ground terms constructible from constants (including all integers) and function symbols by  $\mathcal{T}$ , and let  $\mathcal{C}$  stand for the subset of symbolic (i.e. non-integer) constants. The *ground instance* of  $P$ , denoted by  $grd(P)$ , is the set of all ground rules constructible from rules  $r \in P$  by substituting every variable in  $r$  with some element of  $\mathcal{T}$ . We associate  $P$  with its *positive atom dependency graph*

$$G(P) = (A(grd(P)), \{(a_0, a) \mid r \in grd(P), h(r) = a_0, a \in B(r)^+\})$$

and call a maximal non-empty subset of  $A(grd(P))$  inducing a strongly connected subgraph<sup>3</sup> of  $G(P)$  a strongly connected component of  $P$ .

A set  $X$  of ground atoms is a *model* of  $P$ , if  $h(r) \in X$ ,  $B(r)^+ \not\subseteq X$ , or  $B(r)^- \cap X \neq \emptyset$  holds for every  $r \in grd(P)$ . Moreover,  $X$  is a *stable model* of  $P$  (Gelfond and Lifschitz 1988), if  $X$  is a  $\subseteq$ -minimal model of  $\{h(r) \leftarrow B(r)^+ \mid r \in grd(P), B(r)^- \cap X = \emptyset\}$ .

Following Oikarinen and Janhunen (2006), a *module*  $\mathbb{P}$  is a triple  $(P, I, O)$  consisting of a ground logic program  $P$  along with sets  $I$  and  $O$  of ground *input* and *output* atoms such that

- (1)  $I \cap O = \emptyset$ ,
- (2)  $A(P) \subseteq I \cup O$ , and
- (3)  $H(P) \subseteq O$ .

We also denote the constituents of  $\mathbb{P} = (P, I, O)$  by  $P(\mathbb{P}) = P$ ,  $I(\mathbb{P}) = I$ , and  $O(\mathbb{P}) = O$ . A set  $X$  of ground atoms is a *stable model* of a module  $\mathbb{P}$ , if  $X$  is a (standard) stable model of  $P(\mathbb{P}) \cup \{a \leftarrow \mid a \in I(\mathbb{P}) \cap X\}$ .

Two modules  $\mathbb{P}_1$  and  $\mathbb{P}_2$  are *compositional*, if

- (1)  $O(\mathbb{P}_1) \cap O(\mathbb{P}_2) = \emptyset$  and
- (2)  $O(\mathbb{P}_1) \cap C = \emptyset$  or  $O(\mathbb{P}_2) \cap C = \emptyset$

for every strongly connected component  $C$  of  $P(\mathbb{P}_1) \cup P(\mathbb{P}_2)$ .

In other words, all rules defining an atom must belong to the same module. And any positive recursion is within modules; no positive recursion is allowed among modules.

Provided that  $\mathbb{P}_1$  and  $\mathbb{P}_2$  are compositional, their *join* is defined as the module

$$\mathbb{P}_1 \sqcup \mathbb{P}_2 = (P(\mathbb{P}_1) \cup P(\mathbb{P}_2), (I(\mathbb{P}_1) \setminus O(\mathbb{P}_2)) \cup (I(\mathbb{P}_2) \setminus O(\mathbb{P}_1)), O(\mathbb{P}_1) \cup O(\mathbb{P}_2)) .$$

The module theorem (Oikarinen and Janhunen 2006) shows that a set  $X$  of ground atoms is a stable model of  $\mathbb{P}_1 \sqcup \mathbb{P}_2$  iff  $X = X_1 \cup X_2$  for stable models  $X_1$  and  $X_2$

<sup>3</sup> That is, each pair of atoms in the subgraph is connected by a path.

of  $\mathbb{P}_1$  and  $\mathbb{P}_2$ , respectively, such that  $X_1 \cap (I(\mathbb{P}_2) \cup O(\mathbb{P}_2)) = X_2 \cap (I(\mathbb{P}_1) \cup O(\mathbb{P}_1))$ .<sup>4</sup> For example, the modules  $\mathbb{P}_1 = (\{a \leftarrow \sim c; c \leftarrow \sim b\}, \{b\}, \{a, c\})$  and  $\mathbb{P}_2 = (\{b \leftarrow a\}, \{a\}, \{b\})$  are compositional, and combining their stable models,  $\{a, b\}$  and  $\{c\}$  for  $\mathbb{P}_1$  as well as  $\{a, b\}$  and  $\emptyset$  for  $\mathbb{P}_2$ , yields the stable models  $\{a, b\}$  and  $\{c\}$  of  $\mathbb{P}_1 \sqcup \mathbb{P}_2 = (P(\mathbb{P}_1) \cup P(\mathbb{P}_2), \emptyset, \{a, b, c\})$ . Unlike that,  $\mathbb{P}'_1 = (\{a \leftarrow b; c \leftarrow \sim a\}, \{b\}, \{a, c\})$  and  $\mathbb{P}_2$  are not compositional because the strongly connected component  $\{a, b\}$  of  $P(\mathbb{P}'_1) \cup P(\mathbb{P}_2)$  includes  $a \in O(\mathbb{P}'_1)$  and  $b \in O(\mathbb{P}_2)$ . Moreover,  $\{a, b\}$  is a stable model of  $\mathbb{P}'_1$  and  $\mathbb{P}_2$  but not of  $P(\mathbb{P}'_1) \cup P(\mathbb{P}_2)$ .

An assignment  $v$  over a set  $A$  of ground atoms is a function  $v : A \rightarrow \{t, f, u\}$  whose range consists of truth values, standing for true, false, and undefined. Given an assignment  $v$ , we define the sets  $V^x = \{a \in A \mid v(a) = x\}$  for  $x \in \{t, f, u\}$ . In what follows, we represent partial assignments like  $v$  either by  $(V^t, V^f)$  or  $(V^t, V^u)$  by leaving the respective variables with default values implicit.

Finally, we use typewriter font and symbols `:-` and `not` instead of  $\leftarrow$  and  $\sim$ , respectively, whenever we deal with source code accepted by *clingo*. In such a case, we also make use of extended language constructs like integrity or cardinality constraints, all of which can be reduced to normal logic programs, as detailed in the literature (cf. Simons *et al.* (2002)).

### 3 Multi-shot solving with *clingo* at a glance

Let us begin with an informal overview of the central features and corresponding language constructs of *clingo*'s multi-shot solving capacities.

A key feature, distinguishing *clingo* from its predecessors, is the possibility to structure (non-ground) input rules into subprograms. To this end, a program can be partitioned into several subprograms by means of the directive `#program`; it comes with a name and an optional list of parameters. Once given in the input, the directive gathers all rules up to the next such directive (or the end of file) within a subprogram identified by the supplied name and parameter list. As an example, two subprograms `base` and `acid(k)` can be specified as follows:

```

1 a(1).
2 #program acid(k).
3 b(k).
4 c(X,k) :- a(X).
5 #program base.
6 a(2).
```

Listing 1. Logic program with `#program` declarations

Note that `base` is a dedicated subprogram (with an empty parameter list): in addition to the rules in its scope, it gathers all rules not preceded by any `#program` directive. Hence, in the above example, the `base` subprogram includes the facts

<sup>4</sup> Note that the module theorem is strictly stronger than the splitting set theorem (Lifschitz and Turner 1994). For instance, there is no non-trivial splitting set of  $\mathbb{P}_1 \sqcup \mathbb{P}_2$ , since neither  $\{a\}$  nor  $\{b\}$  is one.

a(1) and a(2), although, only the latter is in the actual scope of the directive in Line 5. Without further control instructions (see below), *clingo* grounds and solves the `base` subprogram only, essentially, yielding the standard behavior of ASP systems. The processing of other subprograms such as `acid(k)` is subject to explicitly given control instructions.

For such customized control over grounding and solving, a main routine (taking a control object representing the state of *clingo* as argument, here `prg`) can be supplied. For illustration, let us consider two Python main routines<sup>5</sup>:

```

7  #script (python)                7  #script (python)
8  def main(prg):                  8  def main(prg):
9    prg.ground(["base", []])      9    prg.ground(["acid", [42]])
10   prg.solve()                  10   prg.solve()
11 #end.                          11 #end.

```

While the control program on the left matches the default behavior of *clingo*, the one on the right ignores all rules in the `base` program but rather contains a `ground` instruction for `acid(k)` in Line 8, where the parameter `k` is to be instantiated with the term `42`. Accordingly, the schematic fact `b(k)` is turned into `b(42)`, no ground rule is obtained from “`c(X,k) :- a(X)`” due to lacking instances of `a(X)`, and the `solve` command in Line 10 yields a stable model consisting of `b(42)` only. Note that `ground` instructions apply to the subprograms given as arguments, while `solve` triggers reasoning w.r.t. all accumulated ground rules.

In order to accomplish more elaborate reasoning processes, like those of *iclingo* and *oclingo* or other customized ones, it is indispensable to activate or deactivate ground rules on demand. For instance, former initial or goal state conditions need to be relaxed or completely replaced when modifying a planning problem, e.g., by extending its horizon. While the predecessors of *clingo* relied on the `#volatile` directive to provide a rigid mechanism for the expiration of transient rules, *clingo* captures the respective functionalities and customizations thereof in terms of the `#external` directive. The latter goes back to *lparse* (Syrjänen 2001) and was also supported by *clingo*'s predecessors to exempt (input) atoms from simplifications (and fixing them to false). As detailed in the following, the `#external` directive of *clingo* provides a generalization that, in particular, allows for a flexible handling of yet undefined atoms.

For continuously assembling ground rules evolving at different stages of a reasoning process, `#external` directives declare atoms that may still be defined by rules added later on. In terms of module theory, such atoms correspond to inputs, which (unlike undefined output atoms) must not be simplified. For declaring input atoms, *clingo* supports schematic `#external` directives that are instantiated along with the rules of their respective subprograms. To this end, a directive like

```
#external p(X,Y) : q(X,Z), r(Z,Y).
```

<sup>5</sup> The `ground` routine takes a list of pairs as argument. Each such pair consists of a subprogram name (e.g. `base` or `acid`) and a list of actual parameters (e.g. `[]` or `[42]`).

is treated similar to a rule “ $p(X, Y) :- q(X, Z), r(Z, Y)$ ” during grounding. However, the head atoms of the resulting ground instances are merely collected as inputs, whereas the ground rules as such are discarded.

Once grounded, the truth value of external atoms can be changed via the *clingo* API (until the atoms becomes defined by corresponding rules). By default, the initial truth value of external atoms is set to false. For example, with *clingo*'s Python API, `assign_external(self, p(a,b), True)`<sup>6</sup> can be used to set the truth value of the external atom  $p(a,b)$  to true. Among others, this can be used to activate and deactivate rules in logic programs. For instance, the integrity constraint “ $:- q(a,c), r(c,b), p(a,b)$ ” is ineffective whenever  $p(a,b)$  is false.

## 4 Multi-shot solving

Having set the practical stage in the previous section, let us now turn to posing the formal foundations of multi-shot ASP solving. We begin with a characterization of grounding subprograms with external directives in the context of previously grounded subprograms. This provides us with a formal account of the interplay of ground routines with `#program` and `#external` directives. Next, we show how module theory can be used for characterizing the composition of ground subprograms during multi-shot solving. This gives us a precise idea on the successive logic programs contained in the ASP solver at each invocation of `solve`. Finally, all this culminates in an operational semantics for multi-shot solving in terms of state-changing operations.

The concepts introduced in this section are mainly illustrated by succinct, technical examples. More illustration is provided in the next section discussing several use cases in detail.

### 4.1 Parameterizable subprograms

A *program declaration* is of form

$$\#program\ n(p_1, \dots, p_k) \tag{1}$$

where  $n, p_1, \dots, p_k$  are symbolic constants. We call  $n$  the name of the declaration and  $p_1, \dots, p_k$  its parameters. For simplicity, we suppose that different occurrences of program declarations with the same name also share the same parameters (although this is not required by *clingo*). In this way, each name is associated with a unique parameter specification.

The *scope* of a program declaration in a list of rules and declarations consists of the set of all rules and non-program declarations following the directive up to the next program declaration or the end of the list.<sup>7</sup> In Listing 1, the scope of

<sup>6</sup> In order to construct atoms, symbolic terms, or function terms, respectively, the *clingo* API function `Function` has to be used. Hence, the expression  $p(a,b)$  actually stands for `Function("p", [Function("a"), Function("b")])`.

<sup>7</sup> That is, the end of file in practice.

the declaration in Line 2 consists of  $b(k)$  and “ $c(X, k) :- a(X)$ ,” while that in Line 5 contains  $a(2)$ . Given a list  $R$  of (non-ground) rules and declarations along with a non-integer constant  $n$ , we define  $R(n)$  as the set of all (non-ground) rules and (non-program) declarations in the scope of all occurrences of program declarations with name  $n$ . We often refer to  $R(n)$  as a *subprogram* of  $R$ . All rules and non-program declarations outside the scope of any (explicit) program declaration are thought of being implicitly preceded by a “#program base” declaration. Hence, if  $R$  consists of Lines 1–6 above, we get<sup>8</sup>  $R(\text{base}) = \{a(1) \leftarrow, a(2) \leftarrow\}$  and  $R(\text{acid}) = \{b(k) \leftarrow, c(X, k) \leftarrow a(X)\}$ . Each such list  $R$  induces a collection  $(R(c))_{c \in \mathcal{C}}$  of (non-disjoint) subprograms (most of which are empty). For example, all subprograms obtained from Lines 1–6 are empty, except for *base* and *acid*.

Given a name  $n$  with associated parameters  $p_1, \dots, p_k$ , the instantiation of subprogram  $R(n)$  with terms  $t_1, \dots, t_k$  results in the set  $R(n)[p_1/t_1, \dots, p_k/t_k]$ , obtained by replacing in  $R(n)$  each occurrence of  $p_i$  by  $t_i$  for  $1 \leq i \leq k$ .<sup>9</sup> For instance,  $R(\text{acid})[k/42]$  consists of  $b(42)$  and “ $c(X, 42) :- a(X)$ .”

## 4.2 Contextual grounding

The definition of a program’s ground instance  $\text{grd}(P)$  depends on  $P$  and its underlying set of terms. For instance, grounding accordingly program  $R(\text{base}) \cup R(\text{acid})[k/42]$  yields

$$\{a(1) \leftarrow, a(2) \leftarrow\} \cup \{b(42) \leftarrow\} \cup \left\{ \begin{array}{l} c(1, 42) \leftarrow a(1), \\ c(2, 42) \leftarrow a(2), \\ c(42, 42) \leftarrow a(42) \end{array} \right\} \quad (2)$$

In practice,<sup>10</sup> however, rules are grounded relative to a set of atoms, which we refer to as an *atom base*. In our example, this avoids the generation of the irrelevant rule “ $c(42, 42) \leftarrow a(42)$ .” To see this, note that the relevance of instances of “ $c(X, 42) \leftarrow a(X)$ ” depends upon the available ground atoms of predicate  $a/1$ . Now, grounding first  $R(\text{base})$  establishes — in addition to  $\text{grd}(R(\text{base}))$  — the atom base  $\{a(1), a(2)\}$ . Grounding then  $R(\text{acid})[k/42]$  relative to this atom base, yields

$$\{b(42) \leftarrow\} \cup \left\{ \begin{array}{l} c(1, 42) \leftarrow a(1), \\ c(2, 42) \leftarrow a(2) \end{array} \right\} \quad (3)$$

because only rule instances are created if their positive body literals either belong to the atom base or are derivable through other rules instances.<sup>11</sup> Hence, rule

<sup>8</sup> We drop the typewriter font, whenever our emphasis shifts to a more formal context.

<sup>9</sup> *clingo* uses a more general instantiation process involving unification and arithmetic evaluation; see Gebser et al. (2015) for details.

<sup>10</sup> In one-shot grounding, a program is partitioned via the strongly connected components of its dependency graph.

<sup>11</sup> This is a simplification of semi-naive database evaluation (Abiteboul et al. 1995), used in ASP grounding components (Kaufmann et al. 2016). Notably, this technique allows for dealing with recursive function symbols and guarantees termination for a wide class of programs,  $P$ , even though their ground instantiation  $\text{grd}(P)$  is infinite.



$c(42, 42) \leftarrow a(42)$  is dropped. This is made precise in view of our purpose in the following definition.

Given a set  $R$  of (non-ground) rules and two sets  $C, D$  of ground atoms, we define an instantiation of  $R$  relative to atom base  $C$  as a ground program  $grd_C(R)$  over atom base  $D$  subject to the following conditions:

$$D = C \cup H(grd_C(R)) \quad (4)$$

$$grd_C(R) \subseteq \{h(r) \leftarrow B(r)^+ \cup \{\sim a \mid a \in B(r)^- \cap D\} \mid r \in grd(R), B(r)^+ \cup \{h(r)\} \subseteq D\} \quad (5)$$

$$grd_C(R) \cup Q \text{ and } grd(R) \cup Q \text{ have the same stable models} \quad (6)$$

where  $Q = \{\{a\} \leftarrow \mid a \in C \setminus H(grd(R))\}$ .<sup>12</sup>

Atom base  $D$  gives the extension of  $C$  obtained by grounding  $R$  relative to  $C$ . To this end, Condition (4) limits  $D$  to atoms either belonging to  $C$  or emerging as heads of rules in  $grd_C(R)$ , while Condition (5) projects  $grd(P)$  to  $D$  by reducing its rules to those parts possibly relevant to stable models, as expressed in (6). The exact scope of the obtained atom base  $D$  as well as  $grd_C(R)$  are left open to account for potential simplifications during grounding.<sup>13</sup> The correctness of the specific scope is warranted by Condition (6)

We capture the composition of ground programs below in Section 4.4 in terms of modules. To this end, note that the choices in  $Q$  mimic the role of input atoms of modules. In fact, Condition (6) is equivalent to requiring that the modules  $(grd_C(R), C \setminus H(grd(R)), H(grd(R)))$  and  $(grd(R), C \setminus H(grd(R)), H(grd(R)))$  have the same stable models.

Resuming our example, we see that  $grd_{\{a(1), a(2)\}}(R(\text{acid})[k/42])$  corresponds to the rules in (3). Together with  $grd_{\emptyset}(R(\text{base}))$ , the resulting program is equivalent to  $grd(R(\text{base}) \cup R(\text{acid})[k/42])$ , viz. the rules in (2).

For further illustration, consider  $R = \{a(X) \leftarrow f(X), e(X); b(X) \leftarrow f(X), \sim e(X)\}$  along with  $C = \{f(1), f(2), e(1)\}$ . Focusing on relevant (parts of) rule instances relative to  $C$  leads to

$$grd_C(R) = \left\{ \begin{array}{l} a(1) \leftarrow f(1), e(1) \quad b(1) \leftarrow f(1), \sim e(1) \\ \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad b(2) \leftarrow f(2) \end{array} \right\}$$

over  $D = C \cup \{a(1), b(1), b(2)\}$ . In particular, note that an inapplicable rule instance including  $e(2)$  is dropped, while  $\sim e(2)$  is simplified away to thus obtain the last of the above ground rules.

Although  $grd_C(R)$  reflects a potential simplification of the full ground program  $grd(R)$ , both can likewise be augmented with additional rules. Moreover, the restriction of the resulting atom base preserves equivalence. The next result makes this precise in terms of module theory.

<sup>12</sup> A *choice rule* (Simons *et al.* 2002) of the form  $\{a\} \leftarrow$  corresponds to (normal) rules  $a \leftarrow \sim a'$  and  $a' \leftarrow \sim a$ , where  $a'$  is a fresh atom.

<sup>13</sup> In fact, the instantiation process of *clingo* iteratively extends an atom base by heads of rules whose positive body atoms are contained in it. Moreover, the finite instantiation of a program like  $\{a(9); b(1); b(X+1) \leftarrow b(X), \sim a(X)\}$  relies on the evaluation of  $\sim a(X)$  during grounding, allowing *clingo* to stop the successive generation of rule instances at  $b(10) \leftarrow b(9), \sim a(9)$ .

*Proposition 1*

Let  $R$  be a set of (non-ground) rules, let  $grd_C(R)$  be an instantiation of  $R$  relative to a set  $C$  of ground atoms, and let  $\mathbb{P}$  be a module.

If  $\mathbb{P}$  and  $(grd(R), C \setminus H(grd(R)), H(grd(R)))$  are compositional, then  $\mathbb{P}$  and  $(grd_C(R), C \setminus H(grd(R)), H(grd(R)))$  are compositional as well, where  $\mathbb{P} \sqcup (grd(R), C \setminus H(grd(R)), H(grd(R)))$  and  $\mathbb{P} \sqcup (grd_C(R), C \setminus H(grd(R)), H(grd(R)))$  have the same stable models.

*Proof*

Assume that  $\mathbb{P}$  and  $(grd(R), C \setminus H(grd(R)), H(grd(R)))$  are compositional. By the construction of  $grd_C(R)$  in (5),  $G(grd_C(R))$  is a subgraph of  $G(grd(R))$ , which implies that  $\mathbb{P}$  and  $(grd_C(R), C \setminus H(grd(R)), H(grd(R)))$  are compositional as well. As  $(grd(R), C \setminus H(grd(R)), H(grd(R)))$  and  $(grd_C(R), C \setminus H(grd(R)), H(grd(R)))$  are equivalent by the condition in (6), the module theorem (Oikarinen and Janhunen 2006) yields that  $\mathbb{P} \sqcup (grd(R), C \setminus H(grd(R)), H(grd(R)))$  and  $\mathbb{P} \sqcup (grd_C(R), C \setminus H(grd(R)), H(grd(R)))$  have the same stable models.  $\square$

More illustration of contextual grounding is given throughout the following sections.

**4.3 Extensible logic programs**

We define a (non-ground) logic program  $R$  as *extensible*, if it contains some (non-ground) *external declaration* of the form

$$\#external\ a : B \tag{7}$$

where  $a$  is an atom and  $B$  a rule body.

As an example, consider the extensible program in Listing 2.

```

1 #external e(X) : f(X), X < 2.
2 f(1..2).
3 a(X) :- f(X), e(X).
4 b(X) :- f(X), not e(X).

```

Listing 2. Extensible logic program

For grounding an external declaration as in (7), we treat it as a rule  $a \leftarrow B, \varepsilon$  where  $\varepsilon$  is a distinguished ground atom marking rules from  $\#external$  declarations. Formally, given an extensible program  $R$ , we define the collection  $Q$  of rules corresponding to  $\#external$  declarations as follows.

$$Q = \{a \leftarrow B, \varepsilon \mid (\#external\ a : B) \in R\}$$

$$R' = \{a \leftarrow B \in R\}$$

With these, the ground instantiation of an extensible logic program  $R$  relative to an atom base  $C$  is defined as a ground logic program  $P$  associated with a set  $E$  of ground atoms, where

$$P = \{r \in grd_{C \cup \{\varepsilon\}}(R' \cup Q) \mid \varepsilon \notin B(r)\} \tag{8}$$

$$E = \{h(r) \mid r \in grd_{C \cup \{\varepsilon\}}(R' \cup Q), \varepsilon \in B(r)\} \tag{9}$$

For simplicity, we refer to  $P$  and  $E$  as a *logic program with externals*, and we drop the reference to  $R$  and  $C$  whenever clear from the context. Note that after grounding, the special atom  $\varepsilon$  appears neither in  $P$  nor  $E$ . In fact,  $P$  is a logic program over  $C \cup E \cup H(P)$ .

Given the set  $E$  of externals,  $P$  can alternatively be defined as  $\text{grd}_{C \cup E}(R')$ .

Grounding the program in Listing 2 relative to an empty atom base yields the below program with a single external atom, viz.  $e(1)$  :

```

1  f(1) . f(2) .
2  a(1) :- f(1), e(1) .
3  b(1) :- f(1), not e(1) .
4  b(2) :- f(2) .

```

Note how externals influence the result of grounding. While occurrences of  $e(1)$  remain untouched, the atom  $e(2)$  is unavailable and thus set to false according to the condition in (5). In practice, even more simplifications are applied during grounding. For instance, in *clingo*, the established truth of  $f(1)$  and  $f(2)$  leads to their removal in the bodies in Lines 2–4.

Logic programs with externals constitute a major building block of multi-shot solving. Hence, before addressing their composition within a more elaborate formal framework, let us provide some semantic underpinnings. The stable models of such programs are defined relative to a truth assignment on the external atoms. For a program  $P$  with externals  $E$ , we define the set  $I = E \setminus H(P)$  as input atoms of  $P$ . That is, input atoms are externals that are not overridden by rules in  $P$ . Then, given  $P$  along with a partial assignment  $V = (V^t, V^u)$  over  $I$ , we define the stable models of  $P$  w.r.t.  $V$  as the ones of  $P \cup (\{a \leftarrow \mid a \in V^t\} \cup \{\{a\} \leftarrow \mid a \in V^u\})$  to capture the extension of  $P$  with respect to a truth assignment to the input atoms in  $I$ . Note that the externals in  $E$  remain implicit in the domain of  $V$ . For instance, the above program  $P$  with externals  $E = \{e(1)\}$  has a stable model including  $a(1)$  but excluding  $b(1)$  w.r.t. assignment  $(\{e(1)\}, \emptyset)$ , and vice versa with  $(\emptyset, \emptyset)$ .

Further examples involving external atoms are given in Listings 3, 7, 10, 11, and 14 below.

#### 4.4 Composing logic programs with externals

The assembly of (ground) subprograms can be characterized by means of module theory. Program states are captured by modules whose input and output atoms provide the respective atom base. Successive grounding instructions result in modules to be joined with the modules of the corresponding program states.

Given an atom base  $C$ , a (non-ground) extensible program  $R$  yields the module

$$\mathbb{R}(C) = (P, (C \cup E) \setminus H(P), H(P)) \quad (10)$$

via the ground program  $P$  with externals  $E$  obtained by grounding  $R$  relative to  $C$ .<sup>14</sup> For example, grounding the extensible program in Listing 2 w.r.t. the empty atom base results in the module in (11).

$$\left( \left\{ \begin{array}{l} f(1) \leftarrow \\ f(2) \leftarrow \\ a(1) \leftarrow f(1), e(1) \\ b(1) \leftarrow f(1), \sim e(1) \\ b(2) \leftarrow f(2) \end{array} \right\}, \{e(1)\}, \left\{ \begin{array}{l} f(1), f(2), \\ a(1), \\ b(1), b(2) \end{array} \right\} \right) \quad (11)$$

Given the induction of modules from extensible programs w.r.t. to atom bases in (10), we define successive program states in the following way.

The initial program state is given by the empty module  $\mathbb{P}_0 = (\emptyset, \emptyset, \emptyset)$ .

The program state succeeding a module  $\mathbb{P}_i$  is captured by the module

$$\mathbb{P}_{i+1} = \mathbb{P}_i \sqcup \mathbb{R}_{i+1}(I(\mathbb{P}_i) \cup O(\mathbb{P}_i)) \quad (12)$$

where  $\mathbb{R}_{i+1}(I(\mathbb{P}_i) \cup O(\mathbb{P}_i))$  gives the result of grounding an extensible program  $R$  relative to the atom base  $I(\mathbb{P}_i) \cup O(\mathbb{P}_i)$  as defined in (10). Note that  $P(\mathbb{R}_{i+1}(I(\mathbb{P}_i) \cup O(\mathbb{P}_i)))$  is a logic program over  $I(\mathbb{P}_{i+1}) \cup O(\mathbb{P}_{i+1})$  with externals  $E$  such that

$$I(\mathbb{P}_{i+1}) \cup O(\mathbb{P}_{i+1}) = (I(\mathbb{P}_i) \cup O(\mathbb{P}_i)) \cup E \cup H(P(\mathbb{R}_{i+1}(I(\mathbb{P}_i) \cup O(\mathbb{P}_i)))).$$

This reflects the atom base of programs with externals discussed after equations (8) and (9). From a practical point of view, the modules  $\mathbb{P}_i$  and  $\mathbb{P}_{i+1}$  represent the program state of *clingo* before and after the  $(i+1)$ -st execution of a ground command in a main routine. And  $\mathbb{R}_{i+1}(I(\mathbb{P}_i) \cup O(\mathbb{P}_i))$  captures the result of the  $(i+1)$ -st ground command applied to an extensible program relative to the atom base provided by  $\mathbb{P}_i$ . Notably, the join leading to  $\mathbb{P}_{i+1}$  can be undefined in case the constituent modules are non-compositional. At system level, compositionality is only partially checked. *clingo*, or more precisely *clasp*, respectively, issues an error message when atoms become redefined but no cycle check over modules is done.

Interestingly, input atoms induced by externals can also be used to incorporate future information. To see this, consider the following rules (extracted from Listing 10):<sup>15</sup>

```

1 #program s(i).
2 #external a(i).

4 { q(i) }.
5 a(i-1) :- q(i).
6 a(i-1) :- a(i).
7         :- a(i), q(i).

```

<sup>14</sup> Note that  $E \setminus H(P)$  consists of atoms stemming from `#external` declarations that have not been “overwritten” by any rules in  $P$ .

<sup>15</sup> Similar to the  $n$ -Queens problem, putting a queen ( $q$ ) on position  $i$  attacks ( $a$ ) positions  $1, \dots, i-1$ , and no other queen may be put there.

These rules give rise to the module  $\mathbf{S}_i(C)$  given as follows.

$$\left( \left( \begin{array}{l} \{q(i)\} \leftarrow \\ a(i-1) \leftarrow q(i) \\ a(i-1) \leftarrow a(i) \\ \leftarrow a(i), q(i) \end{array} \right), (C \cup \{a(i)\}) \setminus \{a(i-1), q(i)\}, \{a(i-1), q(i)\} \right) \quad (13)$$

We observe that the input atom  $a(i)$  of  $\mathbf{S}_i(C)$  is defined in  $\mathbf{S}_{i+1}(C)$ . Proceeding as in (12) by letting  $\mathbf{R}$  be  $\mathbf{S}$ , each module  $\mathbf{P}_i$  has a single input atom  $a(i)$  and output atoms  $\{a(0), \dots, a(i-1), q(1), \dots, q(i)\}$ ; it yields  $i+1$  stable models, either the empty one or one of the form  $\{a(0), \dots, a(j-1), q(j)\}$  for  $1 \leq j \leq i$ . In the latter models, the truth of an atom like  $a(0)$  relies on that of  $q(j)$ , occurring in a subsequently joined module whenever  $j \geq 2$ .

The next result provides a characterization of sequences of program states in terms of their constituent subprograms and their associated external atoms. The well-definedness of a sequence depends upon their compositionality, as defined in Section 2.

### Proposition 2

Let  $(R_i)_{i \geq 0}$  be a sequence of (non-ground) extensible programs, and let  $P_{i+1}$  be the ground program with externals  $E_{i+1}$  obtained from  $R_{i+1}$  and  $I(\mathbf{P}_i) \cup O(\mathbf{P}_i)$  for  $i \geq 0$ , where  $\mathbf{P}_0 = (\emptyset, \emptyset, \emptyset)$  and  $\mathbf{P}_{i+1}$  is defined as in (12), and  $\mathbf{R}_{i+1}$  is defined as in (10) for  $i \geq 0$ .

If  $\mathbf{P}_i$  and  $\mathbf{R}_{i+1}(I(\mathbf{P}_i) \cup O(\mathbf{P}_i))$  are compositional for some  $j \geq 0$  and all  $j > i \geq 0$ , then

- (1)  $P(\mathbf{P}_j) = \bigcup_{i > 0}^j P_i$
- (2)  $I(\mathbf{P}_j) = \bigcup_{i > 0}^j E_i \setminus \bigcup_{i > 0}^j H(P_i)$
- (3)  $O(\mathbf{P}_j) = \bigcup_{i > 0}^j H(P_i)$

### Proof

For  $\mathbf{P}_0 = (\emptyset, \emptyset, \emptyset)$ , we have that  $P(\mathbf{P}_0) = I(\mathbf{P}_0) = O(\mathbf{P}_0) = \emptyset$ , and assume that  $P(\mathbf{P}_j) = \bigcup_{i > 0}^j P_i$ ,  $I(\mathbf{P}_j) = \bigcup_{i > 0}^j E_i \setminus \bigcup_{i > 0}^j H(P_i)$ , and  $O(\mathbf{P}_j) = \bigcup_{i > 0}^j H(P_i)$  for some  $j \geq 0$ . According to (8)–(10) and (12),  $\mathbf{R}_{j+1}(I(\mathbf{P}_j) \cup O(\mathbf{P}_j))$  is the module  $(P_{j+1}, (I(\mathbf{P}_j) \cup O(\mathbf{P}_j) \cup E_{j+1}) \setminus H(P_{j+1}), H(P_{j+1}))$ , where

$$\begin{aligned} P_{j+1} &= \{r \in \text{grd}_{I(\mathbf{P}_j) \cup O(\mathbf{P}_j) \cup \{\varepsilon\}}(\{a \leftarrow B \in R_{j+1}\} \cup \\ &\quad \{a \leftarrow B, \varepsilon \mid (\# \text{external } a : B) \in R_{j+1}\}) \mid \varepsilon \notin B(r)\} \\ E_{j+1} &= \{h(r) \mid r \in \text{grd}_{I(\mathbf{P}_j) \cup O(\mathbf{P}_j) \cup \{\varepsilon\}}(\{a \leftarrow B \in R_{j+1}\} \cup \\ &\quad \{a \leftarrow B, \varepsilon \mid (\# \text{external } a : B) \in R_{j+1}\}), \varepsilon \in B(r)\} \end{aligned}$$

Provided that  $\mathbf{P}_j$  and  $\mathbf{R}_{j+1}(I(\mathbf{P}_j) \cup O(\mathbf{P}_j))$  are compositional, their join is defined as

$$\mathbf{P}_{j+1} = \left( \begin{array}{l} \bigcup_{i > 0}^{j+1} P_i, \\ \bigcup_{i > 0}^j E_i \setminus \bigcup_{i > 0}^{j+1} H(P_i) \cup (\bigcup_{i > 0}^{j+1} E_i \cup \bigcup_{i > 0}^j H(P_i)) \setminus \bigcup_{i > 0}^{j+1} H(P_i), \\ \bigcup_{i > 0}^{j+1} H(P_i) \end{array} \right)$$

That is,  $P(\mathbb{P}_{j+1}) = \bigcup_{i>0}^{j+1} P_i$ ,  $I(\mathbb{P}_{j+1}) = \bigcup_{i>0}^{j+1} E_i \setminus \bigcup_{i>0}^{j+1} H(P_i)$ , and  $O(\mathbb{P}_{j+1}) = \bigcup_{i>0}^{j+1} H(P_i)$ . □

The ground rules in  $P(\mathbb{R}_{i+1}(I(\mathbb{P}_i) \cup O(\mathbb{P}_i)))$  are obtained by grounding  $R_{i+1}$  relative to the atom base  $I(\mathbb{P}_i) \cup O(\mathbb{P}_i)$ , viz. the previously gathered input and output atoms. Unlike this, the full ground program  $grd(R_{i+1})$  takes all ground terms into account; this includes all integers. Thus, for example, grounding the extensible program in Listing 2 in full yields an infinite module:

$$\left( \left( \begin{array}{l} f(1) \leftarrow \\ f(2) \leftarrow \\ a(1) \leftarrow f(1), e(1) \\ b(1) \leftarrow f(1), \sim e(1) \\ a(2) \leftarrow f(2), e(2) \\ b(2) \leftarrow f(2), \sim e(2) \\ a(3) \leftarrow f(3), e(3) \\ b(3) \leftarrow f(3), \sim e(3) \\ \dots \end{array} \right) , \{e(1)\}, \left\{ \begin{array}{l} f(1), f(2), \\ a(1), a(2), a(3), \dots \\ b(1), b(2), b(3), \dots \end{array} \right\} \right) \quad (14)$$

Both this module and the one in (11) obtained by contextual grounding possess two stable models:  $X_1 = \{e(1), f(1), f(2), a(1), b(2)\}$  and  $X_2 = \{f(1), f(2), b(1), b(2)\}$ . However, when joined with the module  $(\{e(2) \leftarrow\}, \emptyset, \{e(2)\})$ , the stable models of (11) turn into  $X_1 \cup \{e(2)\}$  and  $X_2 \cup \{e(2)\}$ , while (14) yields  $X_1 \setminus \{b(2)\} \cup \{e(2), a(2)\}$  and  $X_2 \setminus \{b(2)\} \cup \{e(2), a(2)\}$ . This mismatch results from the fact that the atom  $e(2)$  was removed from (11) when grounding relative to atom base  $\{e(1)\}$ . The subsequent definition of  $e(2)$  in module  $(\{e(2) \leftarrow\}, \emptyset, \{e(2)\})$  is thus stripped of any logical relation to the rules in (11). Such differences can be eliminated by stipulating  $H(grd(R_{j+1})) \cap \bigcup_{i>0}^j A(grd(R_i)) \subseteq I(\mathbb{P}_j)$  for a sequence  $(R_i)_{i>0}$  of extensible programs and all  $j \geq 0$ , where the program state  $\mathbb{P}_j$  is obtained through contextual grounding. This condition rules out the join conducted in our example. But even so, full and contextual grounding would still be imbalanced. For example, the module in (11) can be joined with modules defining  $a(3), b(3), \dots$ , while doing the same with (14) violates compositionality. Not to mention that infinite ground programs as in (14) cannot be utilized in practice. This discussion shows the influence of contextual grounding on inputs, outputs, and resulting ground programs. Hence, some care has to be taken when writing interacting subprograms. Actually, apart from the ones in Section 5.2, all following modules obtained by grounding parametrized programs satisfy the above condition and have the same solutions no matter which form of grounding is used.

### 4.5 State-based characterization of multi-shot solving

For capturing multi-shot solving, we must account for sequences of system states, involving information about the programs kept within the grounder and the solver. To this end, we define a simple operational semantics based on system states and associated operations.

An ASP system state is a triple  $(\mathbf{R}, \mathbb{P}, V)$  where

- $\mathbf{R} = (R_c)_{c \in \mathcal{C}}$  is a collection of extensible (non-ground) logic programs,<sup>16</sup>
- $\mathbb{P}$  is a module,
- $V = (V^t, V^u)$  is a three-valued assignment over  $I(\mathbb{P})$ .

When solving with  $\mathbb{P}$ , the input atoms in  $I(\mathbb{P})$  are taken to be false by default, that is,  $V^f = I(\mathbb{P}) \setminus (V^t \cup V^u)$ . This can still be altered by dedicated directives as illustrated below.

As informal examples for ASP system states, consider the ones obtained from the program in Listing 1 after separately grounding subprograms `base` and `acid` (while replacing `k` with `42`), respectively:

$$((R(\text{base}), R(\text{acid})), (\{a(1) \leftarrow, a(2) \leftarrow\}, \emptyset, \{a(1), a(2)\}), (\emptyset, \emptyset)) \quad (15)$$

$$((R(\text{base}), R(\text{acid})), (\{b(42) \leftarrow\}, \emptyset, \{b(42)\}), (\emptyset, \emptyset)) \quad (16)$$

Given that the program in Listing 1 has no external declarations, no truth values can be assigned. This is different in states obtained from the program in Listing 2. Grounding this yields the state

$$((R(\text{base})), \mathbb{R}_b, (\emptyset, \emptyset)) \quad (17)$$

where  $\mathbb{R}_b$  is the module given in (11). Furthermore, we have set the external atom  $e(1)$  to false. The way such states are obtained from non-ground logic programs is made precise next.

ASP system states can be created and modified by the following operations.

Function *create* partitions a program into subprograms.

$create(R) : \mapsto (\mathbf{R}, \mathbb{P}, V)$

for a list  $R$  of (non-ground) rules and declarations where

- $\mathbf{R} = (R(c))_{c \in \mathcal{C}}$
- $\mathbb{P} = (\emptyset, \emptyset, \emptyset)$
- $V = (\emptyset, \emptyset)$

Each subprogram  $R(c)$  gathers all rules and non-program directives in the scope of  $c$ .

The respective subprograms can be extended by function *add* with rules as well as external declarations.

$add(R) : (\mathbf{R}_1, \mathbb{P}, V) \mapsto (\mathbf{R}_2, \mathbb{P}, V)$

for a list  $R$  of (non-ground) rules and declarations where

- $\mathbf{R}_1 = (R_c)_{c \in \mathcal{C}}$  and  $\mathbf{R}_2 = (R_c \cup R(c))_{c \in \mathcal{C}}$

Obviously, we have  $add(R)(create(\emptyset)) = create(R)$ . Note that *add* only affects non-ground subprograms, and thus ignores compositionality issues since they appear on the ground level.

<sup>16</sup> Note that  $R_c$  is merely an indexed set, and thus different from  $R(c)$ .

Function *ground* instantiates the designated subprograms in  $R$  and binds their parameters. The resulting ground programs along with their external atoms are then joined with the ones captured in the current state—provided that they are compositional.

$ground((n, \mathbf{t}_n)_{n \in \mathcal{N}}) : (R, \mathbb{P}_1, V_1) \mapsto (R, \mathbb{P}_2, V_2)$

for a collection  $(n, \mathbf{t}_n)_{n \in \mathcal{N}}$  of pairs of non-integer constants  $\mathcal{N} \subseteq \mathcal{C}$  and term tuples  $\mathbf{t}_n \in \mathcal{T}^{k_n}$  of arity  $k_n$  where

- $\mathbb{P}_2 = \mathbb{P}_1 \sqcup \mathbb{R}(I(\mathbb{P}_1) \cup O(\mathbb{P}_1))$   
and  $\mathbb{R}(I(\mathbb{P}_1) \cup O(\mathbb{P}_1))$  is the module obtained as in (10) from
  - extensible program  $\bigcup_{n \in \mathcal{N}} R_n[p_1/t_1, \dots, p_{k_n}/t_{k_n}]$ , where  $\mathbf{t}_n = (t_1, \dots, t_{k_n})$  and
  - atom base  $I(\mathbb{P}_1) \cup O(\mathbb{P}_1)$
- for  $(R_c)_{c \in \mathcal{C}} = R$
- $V_2^t = \{a \in I(\mathbb{P}_2) \mid V_1(a) = t\}$
- $V_2^u = \{a \in I(\mathbb{P}_2) \mid V_1(a) = u\}$

A few more technical remarks are in order. First, note that a previous external status of an atom is eliminated once it becomes defined by a ground rule. This is accomplished by module composition, namely, the elimination of output atoms from input atoms. Second, note that jointly grounded subprograms are treated as a single logic program. In fact, while  $ground((c, \mathbf{p}_c), (c, \mathbf{p}_c))(s)$  and  $ground((c, \mathbf{p}_c))(s)$  yield the same result,  $ground((c, \mathbf{p}_c))(ground((c, \mathbf{p}_c))(s))$  leads to two non-compositional modules whenever normal rules are contained in  $R_c$ . Finally, note that new inputs stemming from just added external declarations are set to false in view of  $V_2^f = I(\mathbb{P}_2) \setminus (V_2^t \cup V_2^u)$ .

The above functionality lets us now formally characterize the states in (15) and (16). By abbreviating the logic program in Listing 1 with  $R$ , the system state in (15) results from  $ground((base, ())) (create(R))$ , while  $ground((acid, (42))) (create(R))$  yields (16). Moreover, observe the difference between grounding subprogram *acid* before *base* and vice versa. While the ground program comprised in

$$ground((base, ())) (ground((acid, (42))) (create(R)))$$

only consists of the facts  $\{b(42) \leftarrow, a(1) \leftarrow, a(2) \leftarrow\}$ , the one contained in

$$ground((acid, (42))) (ground((base, ())) (create(R)))$$

includes additionally the rules  $\{c(1, 42) \leftarrow a(1); c(2, 42) \leftarrow a(2)\}$ . This difference is due to contextual grounding (cf. Section 4.2). While in the first case rule  $c(X, 42) \leftarrow a(X)$  is grounded w.r.t. atom base  $\{b(42)\}$ , it is grounded relative to  $\{a(1), a(2), b(42)\}$  in the second case. Such effects are obviously avoided when jointly grounding both subprograms, as in

$$ground((base, ()), (acid, (42))) (create(R)) .$$

The next function allows us to change the truth assignment of input atoms.

$assignExternal(a, v) : (R, \mathbb{P}, V_1) \mapsto (R, \mathbb{P}, V_2)$

for a ground atom  $a$  and  $v \in \{t, u, f\}$  where



- if  $v = t$ 
  - $V_2^t = V_1^t \cup \{a\}$  if  $a \in I(\mathbb{P})$ , and  $V_2^t = V_1^t$  otherwise
  - $V_2^u = V_1^u \setminus \{a\}$
- if  $v = u$ 
  - $V_2^t = V_1^t \setminus \{a\}$
  - $V_2^u = V_1^u \cup \{a\}$  if  $a \in I(\mathbb{P})$ , and  $V_2^u = V_1^u$  otherwise
- if  $v = f$ 
  - $V_2^t = V_1^t \setminus \{a\}$
  - $V_2^u = V_1^u \setminus \{a\}$

While the default truth value of input atoms is false, making them undefined results in a choice. Note that *assignExternal* only affects input atoms, that is, “non-overwritten” external atoms. If an atom is not external, then *assignExternal* has no effect.

With this function, we can now characterize the system state in (17). Abbreviating the program in Listing 2 with  $R$ , system state (17) is issued by

$$\text{assignExternal}(e(1), f)(\text{ground}((\text{base}, ())) (\text{create}(R))) . \quad (18)$$

The resulting state is the same as the previous one, obtained from  $\text{ground}((\text{base}, ())) (\text{create}(R))$ , since external atoms are assigned false by function *ground*.

Function *releaseExternal* removes the external status from an atom and sets it permanently to false, otherwise this function has no effect.

$$\text{releaseExternal}(a) : (\mathbf{R}, \mathbb{P}_1, V_1) \mapsto (\mathbf{R}, \mathbb{P}_2, V_2)$$

for a ground atom  $a$  where

- $\mathbb{P}_2 = (P(\mathbb{P}_1), I(\mathbb{P}_1) \setminus \{a\}, O(\mathbb{P}_1) \cup \{a\})$  if  $a \in I(\mathbb{P}_1)$ , and  $\mathbb{P}_2 = \mathbb{P}_1$  otherwise
- $V_2^t = V_1^t \setminus \{a\}$
- $V_2^u = V_1^u \setminus \{a\}$

Note that *releaseExternal* only affects input atoms; defined atoms remain unaffected. The addition of  $a$  to the output makes sure that it can never be re-defined, neither by a rule nor an external declaration. A released (input) atom is thus permanently set to false, since it is neither defined by any rule nor part of the input atoms, and is also denied both statuses in the future.

The following properties shed some light on the interplay among the previous operations. For an ASP system state  $s$  and  $v, v' \in \{t, f, u\}$ , we have

- (1)  $\text{releaseExternal}(a)(\text{releaseExternal}(a)(s)) = \text{releaseExternal}(a)(s)$
- (2)  $\text{releaseExternal}(a)(\text{assignExternal}(a, v)(s)) = \text{releaseExternal}(a)(s)$
- (3)  $\text{assignExternal}(a, v)(\text{releaseExternal}(a)(s)) = \text{releaseExternal}(a)(s)$
- (4)  $\text{assignExternal}(a, v)(\text{assignExternal}(a, v')(s)) = \text{assignExternal}(a, v)(s)$

Finally, *solve* leaves the system state intact and outputs a possibly filtered set of stable models of the logic program with externals comprised in the current state (cf. Section 4.3). This set is general enough to define all basic reasoning modes of ASP.

$solve((A^t, A^f)) : (\mathbf{R}, \mathbb{P}, V) \mapsto (\mathbf{R}, \mathbb{P}, V)$

outputs the set

$$\mathcal{X}_{\mathbb{P}, V} = \{X \mid X \text{ is a stable model of } P(\mathbb{P}) \text{ w.r.t. } V \text{ such that } A^t \subseteq X \text{ and } A^f \cap X = \emptyset\} \quad (19)$$

To be more precise, a state like  $(\mathbf{R}, \mathbb{P}, V)$  comprises the ground logic program  $P(\mathbb{P})$  with external atoms  $I(\mathbb{P})$ . The latter constitutes the domain of the partial assignment  $V$ . Recall from Section 4.3 that the stable models of  $P(\mathbb{P})$  w.r.t.  $V$  are given by the stable models of the program  $P(\mathbb{P}) \cup \{a \leftarrow \mid a \in V^t\} \cup \{\{a\} \leftarrow \mid a \in V^u\}$ . In addition to the assignment  $V$  on input atoms, we consider another partial assignment  $(A^t, A^f)$  over an arbitrary set of atoms for filtering stable models; they are commonly referred to as assumptions.<sup>17</sup> Note the difference among input atoms and (filtering) assumptions. While a true input atom amounts to a fact, a true assumption acts as an integrity constraint.<sup>18</sup> Thus, a true assumption must not be unfounded, while a true external atom is exempt from this condition. Also, undefined input atoms are regarded as false, while undefined assumptions remain neutral. Finally, at the solver level, input atoms are a transient part of the representation, while assumptions only affect the assignment of a single search process.

For illustration, observe that applying  $solve()$  to the system state in (17) leaves the state unaffected and outputs a single stable model containing  $b(1)$ . Unlike this, no model is obtained from  $solve((\emptyset, \{b(1)\}))((17))$ . For a complement,  $solve(assignExternal(e(1), u)((17)))$  outputs two models, one with  $a(1)$  and another with  $b(1)$ .

From the viewpoint of operational semantics, a multi-shot ASP solving process can be associated with a sequence of operations  $(o_k)_{k \in K}$ , which induce a sequence  $(\mathbf{R}_k, \mathbb{P}_k, V_k)_{k \in K}$  of ASP system states where

- (1)  $o_0 = create(\mathbf{R})$  for some logic program  $R$
- (2)  $(\mathbf{R}_0, \mathbb{P}_0, V_0) = o_0$
- (3)  $(\mathbf{R}_k, \mathbb{P}_k, V_k) = o_k((\mathbf{R}_{k-1}, \mathbb{P}_{k-1}, V_{k-1}))$  for  $k > 0$

Note that only  $o_0$  creates states while all others map states to states.

For capturing the result of multi-shot solving in terms of stable models, we consider the sequence of sets of stable models obtained at each solving step. More precisely, given a sequence of operations and system states as above, a multi-shot solving process can be associated with the sequence  $(\mathcal{X}_{\mathbb{P}_j, V_j})_{j \in K, o_j = solve((A_j^t, A_j^f))}$  of sets of stable models, where  $\mathcal{X}_{\mathbb{P}, V}$  is defined w.r.t.  $(A^t, A^f)$  as in (19).

All of the above state operations have almost literal counterparts in *clingo*'s APIs. For instance, the `Control` class of the Python API for capturing sys-

<sup>17</sup> In *clingo*, or more precisely in *clasp*, such assumptions are the principal parameter to the underlying `solve` function (see below). The term assumption traces back to Eén and Sörensson (2003); it was used in ASP by Gebser et al. (2008).

<sup>18</sup> That is, the difference between “ $a \leftarrow$ ” and “ $\leftarrow \sim a$ .”

tem states provides the methods `__init__`, `add`, `ground`, `assign_external`, `release_external`, and `solve`.<sup>19</sup>

#### 4.6 Example

Let us demonstrate the above apparatus via the authentic *clingo* program in Listing 3.

```

1  #external p(1;2;3) .
2  p(0) :- p(3) .
3  p(0) :- not p(0) .

5  #program succ(n) .
6  #external p(n+3) .
7  p(n) :- p(n+3) .
8  p(n) :- not p(n+1), not p(n+2) .

10 #script (python)
11 from clingo import Function
12 def main(prg):
13     prg.ground([("base", [])])
14     prg.assign_external(Function("p", [3]), True)
15     prg.solve()
16     prg.assign_external(Function("p", [3]), False)
17     prg.solve()
18     prg.ground([("succ", [1]), ("succ", [2])])
19     prg.solve()
20     prg.ground([("succ", [3])])
21     prg.solve()
22 #end.
```

Listing 3. Example with `#external` and `#program` declarations controlled by a main routine in Python (`simple.lp`)

This program consists of two subprograms, viz. `base` and `succ` given in Lines 1–3 and 5–8, respectively. Note that once the rule in Line 3 is internalized no stable models are obtained whenever its body is satisfied. Since we use the main routine in Lines 10–22 within a `#script` environment, an initial *clingo* object is created for us and bound to variable `prg` (cf. Line 12). This amounts to an implicit call of `create(R)`, where  $R$  is the list of (non-ground) rules and declarations in Lines 1–8 in Listing 3.<sup>20</sup>

The initial *clingo* object gathers all rules and external declarations in the scope of the subprograms `base` and `succ`; its state is captured by

$$(R_0, \mathbb{P}_0, V_0) = ((R(\text{base}), R(\text{succ})), (\emptyset, \emptyset, \emptyset), (\emptyset, \emptyset)) .$$

<sup>19</sup> For a complete listing of functions and classes available in *clingo*'s Python API, see <https://potassco.org/clingo/python-api/current/clingo.html>

<sup>20</sup> Further *clingo* objects could be created with `create(∅)`, and then further augmented and manipulated.

where  $R(\text{base})$  and  $R(\text{succ})$  consist of the non-ground rules and external declarations in Lines 1–3 and 5–8, respectively. Empty subprograms are omitted.

The initial program state induces the atom base  $I(\mathbb{P}_0) \cup O(\mathbb{P}_0) = \emptyset$ .

The `ground` instruction in Line 13 takes the extensible logic program  $R(\text{base})$  along with the empty base of atoms and yields the ground program  $P_1$  with externals  $E_1$ , where

$$\begin{aligned} P_1 &= \{p(0) \leftarrow p(3); p(0) \leftarrow \sim p(0)\} \\ E_1 &= \{p(1), p(2), p(3)\}. \end{aligned}$$

This results in the module  $\mathbb{R}_1(\emptyset) = (P_1, E_1, \{p(0)\})$ , whose join with  $\mathbb{P}_0$  yields

$$\mathbb{P}_1 = \mathbb{P}_0 \sqcup \mathbb{R}_1(\emptyset) = (\{p(0) \leftarrow p(3); p(0) \leftarrow \sim p(0)\}, \{p(1), p(2), p(3)\}, \{p(0)\}).$$

We then obtain the system state

$$(\mathbf{R}_1, \mathbb{P}_1, V_1) = (\mathbf{R}_0, \mathbb{P}_0 \sqcup \mathbb{R}_1(\emptyset), V_0).$$

While the input atoms  $p(1)$ ,  $p(2)$ , and  $p(3)$  are assigned (by default) to false by  $V_1$ , the instruction in Line 14 switches the value of  $p(3)$  to true. And, we obtain the system state

$$(\mathbf{R}_2, \mathbb{P}_2, V_2) = (\mathbf{R}_0, \mathbb{P}_1, (\{p(3)\}, \emptyset)).$$

Applying the `solve` instruction in Line 15 leaves the state intact and outputs the stable model  $\{p(0), p(3)\}$  of  $\mathbb{P}_2$  w.r.t.  $V_2$ . Note that making  $p(3)$  true leads to the derivation of  $p(0)$ , which blocks the rule in Line 3.

Next, the instruction in Line 16 turns  $p(3)$  back to false, which puts the ASP system into the state

$$(\mathbf{R}_3, \mathbb{P}_3, V_3) = (\mathbf{R}_0, \mathbb{P}_1, (\emptyset, \emptyset)).$$

The last change withdraws the derivation of  $p(0)$  and no stable model is obtained from  $\mathbb{P}_3$  w.r.t.  $V_3$  in Line 17.

The `ground` instruction in Line 18 instantiates the rules and external declarations of subprogram `succ(n)` in Lines 5–8 twice. Once the parameter `n` is instantiated with 1 and once with 2. This yields the extensible logic program  $R_4 = R(\text{succ})[n/1] \cup R(\text{succ})[n/2]$ . This program is then grounded relative to  $I(\mathbb{P}_3) \cup O(\mathbb{P}_3) = \{p(1), p(2), p(3)\} \cup \{p(0)\}$ , which results in the following ground program with externals and resulting module:

$$\begin{aligned} P_4 &= \left\{ \begin{array}{l} p(1) \leftarrow p(4); p(1) \leftarrow \sim p(2), \sim p(3); \\ p(2) \leftarrow p(5); p(2) \leftarrow \sim p(3), \sim p(4) \end{array} \right\} \\ E_4 &= \{p(4), p(5)\} \end{aligned}$$

$$\text{and } \mathbb{R}_4(I(\mathbb{P}_3) \cup O(\mathbb{P}_3)) = \left( P_4, \left\{ \begin{array}{l} p(0), p(4), \\ p(3), p(5) \end{array} \right\}, \left\{ \begin{array}{l} p(1), \\ p(2) \end{array} \right\} \right)$$

Joining the latter with the program module  $\mathbb{P}_3$  of the previous system state yields  $\mathbb{P}_4 = \mathbb{P}_3 \sqcup \mathbb{R}_4(I(\mathbb{P}_3) \cup O(\mathbb{P}_3))$ , or more precisely:

$$\mathbb{P}_4 = \left( \left\{ \begin{array}{l} p(0) \leftarrow p(3); \quad p(1) \leftarrow p(4); \quad p(1) \leftarrow \sim p(2), \sim p(3); \\ p(0) \leftarrow \sim p(0); \quad p(2) \leftarrow p(5); \quad p(2) \leftarrow \sim p(3), \sim p(4) \end{array} \right\}, \left\{ \begin{array}{l} p(4), \\ p(3), p(5) \end{array} \right\}, \right. \\ \left. \left\{ \begin{array}{l} p(0), p(1), \\ p(2) \end{array} \right\} \right)$$

This puts the ASP system into the state

$$(\mathbf{R}_4, \mathbb{P}_4, V_4) = (\mathbf{R}_0, \mathbb{P}_3 \sqcup \mathbb{R}_4(I(\mathbb{P}_3) \cup O(\mathbb{P}_3)), V_3).$$

The subsequent `solve` command in Line 19 leaves the state intact but returns no stable models for  $\mathbb{P}_4$  w.r.t.  $V_4$ .

Then, *clingo* proceeds in Line 20 with the `ground` instruction instantiating  $R(\text{succ})[n/3]$  relative to the atom base  $I(\mathbb{P}_4) \cup O(\mathbb{P}_4) = \{p(0), p(1), p(2), p(3), p(4), p(5)\}$ . This results in the following ground program with externals and induced module:

$$P_5 = \{p(3) \leftarrow p(6); \quad p(3) \leftarrow \sim p(4), \sim p(5)\} \\ E_5 = \{p(6)\} \\ \text{and } \mathbb{R}_5(I(\mathbb{P}_4) \cup O(\mathbb{P}_4)) = \left( P_5, \left\{ \begin{array}{l} p(0), p(1), p(2), \\ p(4), p(5), p(6) \end{array} \right\}, \{p(3)\} \right)$$

With the latter, we obtain the system state

$$(\mathbf{R}_5, \mathbb{P}_5, V_5) = (\mathbf{R}_0, \mathbb{P}_4 \sqcup \mathbb{R}_5(I(\mathbb{P}_4) \cup O(\mathbb{P}_4)), V_3)$$

where

$$\mathbb{P}_5 = \left( \left\{ \begin{array}{l} p(0) \leftarrow p(3); \quad p(1) \leftarrow p(4); \quad p(1) \leftarrow \sim p(2), \sim p(3); \\ p(0) \leftarrow \sim p(0); \quad p(2) \leftarrow p(5); \quad p(2) \leftarrow \sim p(3), \sim p(4); \\ p(3) \leftarrow p(6); \quad p(3) \leftarrow \sim p(4), \sim p(5) \end{array} \right\}, \left\{ \begin{array}{l} p(4), \\ p(5), \\ p(6) \end{array} \right\}, \right. \\ \left. \left\{ \begin{array}{l} p(0), p(1), \\ p(2), \\ p(3) \end{array} \right\} \right).$$

Finally, the `solve` command in Line 21 yields the stable model  $\{p(0), p(3)\}$  of module  $\mathbb{P}_5$  w.r.t. assignment  $(\emptyset, \emptyset)$ .

The result of the ASP solving process induced by program `simple.lp` from Listing 3 is given in Listing 4. The parameter 0 instructs *clingo* to compute all stable models upon each invocation of `solve`.<sup>21</sup>

```
$ clingo simple.lp 0
clingo version 4.5.4
Reading from simple.lp
Solving...
Answer: 1
p(3) p(0)
Solving...
Solving...
```

<sup>21</sup> In fact, *clingo*'s API allows for changing solver configurations in between successive solver calls.

```

Solving...
Answer: 1
p(3) p(0)
SATISFIABLE

Models      : 2
Calls       : 4
Time        : 0.008s (Solving: 0.00s 1st Model: 0.00s Unsat: 0.00s)
CPU Time    : 0.000s

```

Listing 4. Running the program in Listing 3 with *clingo*

Each such invocation is indicated by “Solving...” We see that stable models are only obtained for the first and last invocation. Semantically, our ASP solving process thus results in a sequence of four sets of stable models, namely  $(\{\{p(0), p(3)\}\}, \emptyset, \emptyset, \{\{p(0), p(3)\}\})$ .

The above example illustrates the customized selection of (non-ground) subprograms to instantiate upon `ground` commands. For a convenient declaration of input atoms from other subprogram instances, schematic `#external` declarations are embedded into the grounding process. Given that they do not contribute ground rules, but merely qualify (undefined) atoms that should be exempted from simplifications, `#external` declarations only contribute to the signature of subprograms’ ground instances. Hence, it is advisable to condition them by domain predicates<sup>22</sup> (Syrjänen 2001) only, as this precludes any interferences between signatures and grounder implementations. As long as input atoms remain undefined, their truth values can be freely picked and modified in-between `solve` commands via `assign_external` instructions. This allows for configuring the inputs to modules representing system states in order to select among their stable models. Unlike that, the predecessors *iclingo* and *oclingo* of *clingo* always assigned input atoms to false, so that the addition of rules was necessary to accomplish switching truth values as in Lines 14 and 16 above. However, for a well-defined semantics, *clingo* like its predecessors builds on the assumption that modules resulting from subprogram instantiation are compositional, which essentially requires definitions of atoms and mutual positive dependencies to be local to evolving ground programs (cf. Gebser et al. (2008)).

## 5 Using multi-shot solving in practice

After fixing the formal foundations of multi-shot solving and sketching the corresponding *clingo* constructs, let us now illustrate their usage in several case studies.

### 5.1 Incremental ASP solving

As mentioned, the new *clingo* series fully supersedes its special-purpose predecessors *iclingo* and *oclingo*. To illustrate this, we give below a Python implementation of

<sup>22</sup> Domain and built-in predicates have unique extensions that can be evaluated entirely by means of grounding.

```

1 #program base.
2 on(D,P,0) :- init_on(D,P).

4 #program step(t).
5 1 { move(D,P,t) : disk(D), peg(P) } 1.

7 move(D,t)          :- move(D,P,t).
8 on(D,P,t)          :- move(D,P,t).
9 on(D,P,t)          :- on(D,P,t-1), not move(D,t).
10 blocked(D-1,P,t) :- on(D,P,t-1).
11 blocked(D-1,P,t) :- blocked(D,P,t), disk(D).

13 :- move(D,P,t), blocked(D-1,P,t).
14 :- move(D,t), on(D,P,t-1), blocked(D,P,t).
15 :- disk(D), not 1 { on(D,P,t) } 1.

17 #program check(t).
18 :- goal_on(D,P), not on(D,P,t), query(t).

20 #show move/3.
Listing 5. Towers of Hanoi incremental encoding (tohE.lp)

```

*iclingo*'s control loop, corresponding to the one shipped with *clingo*.<sup>23</sup> Roughly speaking, *iclingo* offers a step-oriented, incremental approach to ASP that avoids redundancies by gradually processing the extensions to a problem rather than repeatedly re-processing the entire extended problem (as in iterative deepening search). To this end, a program is partitioned into a base part, describing static knowledge independent of the step parameter  $t$ , a cumulative part, capturing knowledge accumulating with increasing  $t$ , and a volatile part specific for each value of  $t$ . These parts are delineated in *iclingo* by the special-purpose directives `#base`, `"#cumulative t,"` and `"#volatile t."` In *clingo*, all three parts are captured by `#program` declarations along with `#external` atoms for handling volatile rules. More precisely, our exemplar relies upon subprograms named `base`, `step`, and `check` along with external atoms of form `query(t)`.

We illustrate this approach by adapting the Towers of Hanoi encoding by Gebser *et al.* (2012) in Listing 5. The problem instance in Listing 6 as well as Line 2 in 5 constitute static knowledge, and thus belong to the base program. The transition function is described in the subprogram `step` in Lines 4–15 of Listing 5. Finally, the query is expressed in Line 18; its volatility is realized by making the actual goal condition `"goal_on(D,P), not on(D,P,t)"` subject to the truth assignment

<sup>23</sup> The source code is also available in *clingo*'s examples. The code for incremental solving is in <https://github.com/potassco/clingo/tree/master/examples/clingo/iclingo> and the Towers of Hanoi example in <https://github.com/potassco/clingo/tree/master/examples/gringo/toh>.

```

1  peg(a;b;c) .
2  disk(1..4) .
3  init_on(1..4,a) .
4  goal_on(1..4,c) .

```

Listing 6. Towers of Hanoi instance (toh1.lp)

```

1  #script (python)

3  from clingo import Function

5  def get(val, default):
6      return val if val != None else default

8  def main(prg):
9      imin = get(prg.get_const("imin"), 1)
10     imax = prg.get_const("imax")
11     istop = get(prg.get_const("istop"), "SAT")

13     step, ret = 0, None
14     while ((imax is None or step < imax) and
15            (step == 0 or step < imin or (
16              (istop == "SAT" and not ret.satisfiable) or
17              (istop == "UNSAT" and not ret.unsatisfiable) or
18              (istop == "UNKNOWN" and not ret.unknown)))):
19         parts = []
20         parts.append(("check", [step]))
21         if step > 0:
22             prg.release_external(Function("query", [step-1]))
23             parts.append(("step", [step]))
24             prg.cleanup()
25         else:
26             parts.append(("base", []))
27         prg.ground(parts)
28         prg.assign_external(Function("query", [step]), True)
29         ret, step = prg.solve(), step+1
30 #end.

32 #program check(t) .
33 #external query(t) .

```

Listing 7. Python script implementing *iclingo* functionality in *clingo* (inc.lp)

to the external atom `query(t)`. For convenience, this atom is predefined in Line 33 in Listing 7 as part of the `check` program (cf. Line 32). Hence, subprogram `check` consists of a user- and pre-defined part. Since the encoding of the Towers of Hanoi problem is fairly standard, we refer the interested reader to the literature (Gebser et al. 2012) and devote ourselves in the sequel to its solution by means of multi-shot solving.

Grounding and solving of the program in Listing 6 and 5 is controlled by the Python script in Listing 7. Lines 5–11 fix the values of the constants `imin`,



`imax`, and `istop`. In fact, the setting in Lines 9 and 11 relieves us from adding “`-c imin=0 -c istop="SAT"`” when calling *clingo*. All three constants mimic command line options in *iclingo*. `imin` and `imax` prescribe a least and largest number of iterations, respectively; `istop` gives a termination criterion. The initial values of variables `step` and `ret` are set in Line 13. The value of `step` is used to instantiate the parametrized subprograms and `ret` comprises the solving result. Together, the previous five variables control the loop in Lines 14–29.

The subprograms grounded at each iteration are accumulated in the list `parts`. Each of its entries is a pair consisting of a subprogram name along with its list of actual parameters. In the very first iteration, the subprograms `base` and `check(0)` are grounded. Note that this involves the declaration of the external atom query(0) and the assignment of its default value `false`. The latter is changed in Line 28 to `true` in order to activate the actual query. The `solve` call in Line 29 then amounts to checking whether the goal situation is already satisfied in the initial state. As well, the value of `step` is incremented to 1.

As long as the termination condition remains unfulfilled, each following iteration takes the respective value of variable `step` to replace the parameter in subprograms `step` and `check` during grounding. In addition, the current external atom query( $\tau$ ) is set to `true`, while the previous one is permanently set to `false`. This disables the corresponding instance of the integrity constraint in Line 18 of Listing 5 before it is replaced in the next iteration. In this way, the query condition only applies to the current horizon.

An interesting feature is given in Line 24. As its name suggests, this function cleans up atom bases used during grounding. That is, whenever the truth value of an atom is ultimately determined by the solver, it is communicated to the grounder where it can be used for simplifications in subsequent grounding steps. The call in Line 24 effectively removes atoms from the current atom base (and marks some atoms as facts, which might lead to further simplifications).

The result of each call to `solve` is printed by *clingo*. In our example, the solver is called 16 times before a plan of length 15 is found:

```
$ clingo inc.lp tohE.lp tohI.lp 0
clingo version 4.5.4
Reading from inc.lp ...
Solving...
[...]
Solving...
Answer: 1
move(4,b,1)  move(3,c,2)  move(4,c,3)  move(2,b,4)  move(4,a,5)  \
move(3,b,6)  move(4,b,7)  move(1,c,8)  move(4,c,9)  move(3,a,10)  \
move(4,a,11) move(2,c,12) move(4,b,13) move(3,c,14) move(4,c,15)
SATISFIABLE

Models      : 1
Calls       : 16
Time        : 0.020s (Solving: 0.00s 1st Model: 0.00s Unsat: 0.00s)
CPU Time    : 0.020s
```

Listing 8. Running the programs in Listing 5 and 6 with *clingo*

step	Operation	Line
	<i>create</i> (TOH)	8
0	<i>ground</i> (( <b>base</b> , ()), ( <b>check</b> , (0)))	27
	<i>assignExternal</i> ( <b>query</b> (0), <b>t</b> )	28
	<i>solve</i> (( $\emptyset$ , $\emptyset$ ))	29
1	<i>releaseExternal</i> ( <b>query</b> (0))	22
	<i>ground</i> (( <b>step</b> , (1)), ( <b>check</b> , (1)))	27
	<i>assignExternal</i> ( <b>query</b> (1), <b>t</b> )	28
	<i>solve</i> (( $\emptyset$ , $\emptyset$ ))	29
	⋮	
$k$	<i>releaseExternal</i> ( <b>query</b> ( $k-1$ ))	22
	<i>ground</i> (( <b>step</b> , ( $k$ )), ( <b>check</b> , ( $k$ )))	27
	<i>assignExternal</i> ( <b>query</b> ( $k$ ), <b>t</b> )	28
	<i>solve</i> (( $\emptyset$ , $\emptyset$ ))	29

Fig. 1. Trace of Listing 8 in terms of operations.

For a complement, we give in Figure 1 a trace of the Python script in terms of the operations defined in Section 4.5. We let TOH stand for the combination of programs `tohI.lp` and `tohE.lp` in Listing 5 and 6. Without setting any constants in Listing 7, the sequence stops at the first  $k \geq 0$  for which *solve*(( $\emptyset$ ,  $\emptyset$ )) yields a stable model. Each  $k$ th invocation of *solve*(( $\emptyset$ ,  $\emptyset$ )) is applied to a system state consisting of

- (1) the non-ground programs  $R(\mathbf{base})$ ,  $R(\mathbf{check})$ , and  $R(\mathbf{step})$ ,
- (2) the module obtained by
  - (a) composing the ground subprograms of `base`, `check(0)`, `check(1)`, and `step(1)` for  $1 \leq l \leq k$ ,
  - having
  - (b) the single input atom `query(k)`, and
  - (c) output atoms stemming from
    - (i) all ground rule heads in the subprograms and
    - (ii) all released variables `query(l)` for  $1 \leq l \leq k$ ,
  - and
- (3) a partial assignment mapping `query(k)` to true.

Note that all released atoms `query(l)` are undefined and set to false under stable models semantics. Hence, among all instances of the integrity constraint in Line 18 in Listing 5, only the  $k$ -th one is effective.

## 5.2 *n*-Queens problem

In this section, we consider the well-known  $n$ -Queens problem. However, in contrast to the classical setting, we aim at solving series of problems of increasing size.

### 5.2.1 Encoding incremental cardinality constraints

The  $n$ -Queens problem can be expressed in terms of cardinality constraints, that is, there is exactly one queen per row and column, and there is at most one queen

per diagonal. Hence, for addressing this problem incrementally, we have to encode such constraints in an incremental way.<sup>24</sup> To this end, let us elaborate our encoding technique in a slightly simpler setting. Let  $1, \dots, n$  be a sequence of adjacent positions, such as a row, column, or diagonal, and let  $q_1, \dots, q_n$  be atoms indicating whether a queen is on position  $1, \dots, n$  of such a sequence, respectively.<sup>25</sup>

We begin with a simple way to encode at-most-one constraints for an increasing set of positions  $n$ . The corresponding program,  $Q_i^{\leq 1}$ , is given in Listing 9.

```

1 #program step(i) .
2 { q(i) }.
3 a(i) :- q(i-1) .
4 a(i) :- a(i-1) .
5 :- a(i), q(i) .

```

Listing 9. Incremental encoding of at-most-one constraints

We use  $q(i)$  to represent  $q_i$  as well as auxiliary variables of form  $a(i)$  to indicate that position  $i$  is attacked by a queen on a position  $j \leq i$ . The idea is to join the instantiation of  $Q_n^{\leq 1}$  with the previous program modules whenever a new position  $n$  is added. With this addition a queen may be put on position  $n$  in Line 2. Position  $n$  is attacked if either the directly adjacent position or another connected position is occupied by a queen (Lines 3 and 4). Finally, a queen must not be placed on an attacked position (Line 5).

Let us make this precise by means of the operations introduced in Section 4.5. At first,  $create(Q_i^{\leq 1})$  yields a state comprising  $R(\text{step})$ , an empty module, and an empty assignment. Applying  $ground((\text{step}, (1)))$  to the resulting state yields the module

$$\mathbb{P}_1 = (\{\{q_1\} \leftarrow\}, \emptyset, \{q_1\})$$

and leaves  $R(\text{step})$  as well as the assignment intact.<sup>26</sup> Note that grounding  $Q_1^{\leq 1}$  relative to the empty atom base produces no instances of the rules in Lines 3–5.

Applying  $ground((\text{step}, (2)))$  to the resulting state yields the module

$$\begin{aligned} \mathbb{P}_2 &= (\{\{q_1\} \leftarrow\}, \emptyset, \{q_1\}) \sqcup \left( \left( \begin{array}{l} \{q_2\} \leftarrow \\ a_2 \leftarrow q_1 \\ \leftarrow a_2, q_2 \end{array} \right), \{q_1\}, \{q_2, a_2\} \right) \\ &= \left( \left( \begin{array}{l} \{q_1\} \leftarrow \\ \{q_2\} \leftarrow \\ a_2 \leftarrow q_1 \\ \leftarrow a_2, q_2 \end{array} \right), \emptyset, \{q_1, q_2, a_2\} \right) \end{aligned}$$

Grounding  $Q_2^{\leq 1}$  relative to the output atoms  $\{q_1\}$  of  $\mathbb{P}_1$  produces no instance of the rule in Line 4.

<sup>24</sup> The source code can also be found in *clingo*'s examples: <https://github.com/potassco/clingo/tree/master/examples/clingo/incqueens>

<sup>25</sup> Such sequences are successively build for rows, columns, and diagonals via predicate `target/6` in Listing 11.

<sup>26</sup> Since neither is changed in the sequel, we concentrate on the evolution of the program module.



Accordingly, the corresponding join  $\mathbb{P}_n = \mathbb{Q}_1^=1 \sqcup \dots \sqcup \mathbb{Q}_n^=1$  comprises the union of the programs in (24), and (25)–(26); it has input  $\{a_n\}$  and outputs  $\{q_1, \dots, q_n\} \cup \{a_1, \dots, a_{n-1}\}$ . Then,  $X$  is a stable model of  $\mathbb{P}_n$  iff  $X = \{a_1, \dots, a_{i-1}, q_i\}$  for some  $1 \leq i \leq n$ . That is, the stable models of  $\mathbb{P}_n$  are in a one-to-one correspondence to one element subsets of  $\{q_1, \dots, q_n\}$ .

Interestingly, the last encoding can be turned into one for an at-most one-constraint by omitting Line 7, and into an at-least-one constraint by removing Line 6. Note that the integrity constraint in Line 7 cannot simply be added to the encoding in Listing 9 because it encodes the attack direction the other way round. One could add “:- not a(i), not q(i), query(i).” subject to the query atom. But this has the disadvantage that the constraint would have to be retracted whenever the query atom becomes permanently false. The encoding in Listing 10 ensures that all constraints in the solver (including learnt constraints) can be reused in successive solving steps.

Finally, note that by using `step`, both encodings can be used with the built-in incremental mode, described in the previous section.

### 5.2.2 An incremental encoding

In what follows, we use the above encoding schemes to model the  $n$ -Queens problem. As mentioned, we aim at solving series of differently sized boards. Given that larger boards subsume smaller ones, an evolving problem specification can reuse ground rules from previous *clingo* states when the size increases. To this end, we view the increment of  $n$  to  $n + 1$  as the addition of one more row and column. The basic idea of our incremental encoding is to interconnect the previous and added board cells so that each of them has a unique predecessor or successor in either of the four attack directions of queens, respectively. Each such connection scheme amounts to a sequence of adjacent positions, as used above in Section 5.2.1. The four schemes obtained in our setting are depicted in Figures 2(a)–(d). Direct links are indicated by arrows to target cells with a (white or black) circle. Paths represent the respective ways of attack across several board extensions. They concretise the sequences discussed in the previous section. Attacks from prospective board positions are indicated by black circles (and implemented as external atoms), the ones from the board by white ones. Figure 2(a) illustrates the scheme for backward diagonals. It connects cells of the uppermost previous row to corresponding attacked cells in a new column; the latter are in turn linked to the new cells they attack in the row above. For instance, position (2, 2) is linked to (3, 1), which is itself linked to (1, 3). Note that this scheme ensures that, starting from the middle of any backward diagonal, all cells that are successively added (and belong to the same backward diagonal) are on a path. Such a path follows the board evolution and is directed from previous to newly added cells, where white circles in arrow targets indicate the presence of attacking cells on the board when their respective target cells are added. This orientation is analogous to the above encoding of at-most-one constraints. The schemes for attacks along forward diagonals, horizontal rows, and vertical columns are shown in Figures 2(b)–(d). Notably, the latter two (partially) link new cells to

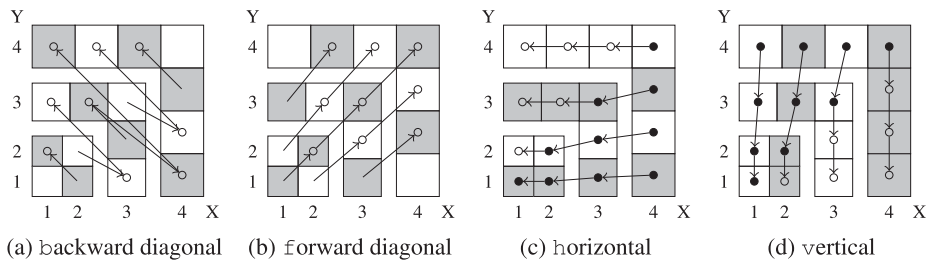


Fig. 2. Attack target links among cells of successive  $n$ -Queens boards up to size 4. (a) Backward diagonal. (b) Forward diagonal. (c) Horizontal. (d) Vertical.

previous ones, in which case the targets are highlighted by black circles. At the level of modules, links from cells that may be added later on give rise to input atoms.

As mentioned, the  $n$ -Queens problem can be expressed in terms of cardinality constraints, requiring that there is exactly one queen per row and column and at most one per diagonal. The idea is to use the incremental encodings from Section 5.2.1 in combination with the four connection schemes depicted in Figures 2(a)–(d). The first encoding, capturing at-most-one constraints, is used for each diagonal, and the second one, handling exactly-one constraints, for each row and column. The resulting *clingo* encoding is given in Listing 11. Let us first outline its structure in relation to the ideas presented in Section 5.2.1. The rules in Lines 7–13 gather linked positions for at-most- and exactly-one constraints in the predicate `target/6`. The choice rule in Line 15 places queens on the new column and row. The rules in Lines 17 and 18 determine which cells are attacked. The rule in Line 20 ensures the at-most-one condition for rows, columns, and diagonals. And finally the rules in Lines 22–23 ensure the at-least-one condition for rows and columns.

Let us make this precise in what follows. After declaring `queen/2` as the output predicate to be displayed, the (sub)program `board(n)` provides rules for extending a board of size  $n-1$  to  $n \geq 1$ . To this end, the `#external` directives in Lines 4 and 5 declare atoms representing horizontal and vertical attacks on cells in the  $n$ th column or row as inputs, respectively. This is analogous to the use of external atoms in Listing 10. Such atoms match the targets of arrows leading to cells with black circles in Figures 2(c) and (d). For instance, `attack(2,1,h)` and `attack(2,2,h)` as well as `attack(1,2,v)` and `attack(2,2,v)` are the inputs to `board(n/2)`. These external atoms express that cells at the horizontal and vertical borders can become targets of attacks once the board is extended beyond size 2. The instances of `target(X,Y,X',Y',D,n)` specified in Lines 7–13 provide links from cells  $(X,Y)$  to targets  $(X',Y')$  along with directions  $D$  leading from or to some newly added cell in the  $n$ th column or row. These instances correspond to arrows shown in Figures 2(a)–(d), yet omitting those to border cells such that `attack(X',Y',D)` is declared as input in Lines 4 and 5, also highlighted by black circles in Figure 2(c) and (d). Queens at newly added cells in the  $n$ th column or row are enabled via the choice rule in Line 15, and the links provided by instances of `target(X,Y,X',Y',D,n)` are utilized in Lines 17 and 18 for deriving `attack(X',Y',D)` in view of a queen at cell  $(X,Y)$  or any of its predecessors in the direction indicated by  $D$ . For instance,

```

1  #show queen/2.

3  #program board(n).
4  #external attack(n,1..n,h).
5  #external attack(1..n,n,v).

7  target(n, X, X, n, b,n) :- X = 1..n-1.
   % diagonal b
8  target(Y, n-1,n, Y-1,b,n) :- Y = 2..n-1.
   % diagonal b
9  target(X, n-1,X+1,n, f,n) :- X = 1..n-1.
   % diagonal f
10 target(n-1,Y, n, Y+1,f,n) :- Y = 1..n-2.
   % diagonal f
11 target(X, n, X-1,n, h,n) :- X = 2..n.
   % horizontal
12 target(n, Y, n-1,Y, h,n) :- Y = 1..n-1.
   % horizontal
13 target(Y, X, Y, X-1,v,n) :- target(X,Y,X-1,Y,h,n). % vertical

15 { queen(1..n,n); queen(n,1..n-1) }.

17 attack(X',Y',D) :- target(X,Y,X',Y',D,n), queen(X,Y).
18 attack(X',Y',D) :- target(X,Y,X',Y',D,n), attack(X,Y,D).

20 :- target(X,Y,X',Y',D,n), attack(X',Y',D), queen(X',Y').

22 :- not queen(1,n), not attack(1,n,h).
23 :- not queen(n,1), not attack(n,1,v).

25 #script (python)
26 def main(prg):
27     n = 0
28     for bound in prg.get_const("calls").arguments:
29         while n < bound.arguments[1].number:
30             n += 1
31             prg.ground(["board", [n]])
32             if n >= bound.arguments[0].number:
33                 print('SIZE {0}'.format(n))
34                 prg.solve()
35 #end.

```

Listing 11. *clingo* program for successive  $n$ -Queens solving (queens.lp)

the following ground rules, simplified by dropping atoms of the domain predicate `target/6`, capture horizontal attacks along the first row of a board of size 4:

```

attack(1,1,h) :- queen(2,1).   attack(1,1,h) :- attack(2,1,h).
attack(2,1,h) :- queen(3,1).   attack(2,1,h) :- attack(3,1,h).
attack(3,1,h) :- queen(4,1).   attack(3,1,h) :- attack(4,1,h).

```

Note that a queen represented by an instance of `queen(X,1)`, for  $2 \leq X \leq 4$ , propagates to cells on its left via an implication chain deriving `attack(X',1,h)`

for every  $1 \leq X' < X$ . Moreover, the fact that the cell at (4,1) can be attacked from the right when increasing the board size is reflected by the input atom `attack(4,1,h)` declared in `board(n/4)`. Given that attacks are propagated analogously for other rows and directions, instances of the integrity constraint in Line 20 prohibit a queen at cell  $(X',Y')$  whenever `attack(X',Y',D)` signals that some predecessor in either direction  $D$  has a queen already. The integrity constraints in Lines 22 and 23 additionally require that each row and column contains some queen. In view of the orientations of horizontal and vertical links, as displayed in Figures 2(c) and (d), non-emptiness can be recognized from a queen at or an attack propagated to the first position in a row or column, no matter to which size the board is extended later on. The incremental development of these rules is illustrated in Figure 3 for the first row of a board (abbreviating predicates by their first letter). It is interesting to observe that the rules generated at each step in Lines 15–22 correspond to the ones in (24) to (26).

Importantly, instantiations of `board(n)` with different integers for  $n$  define distinct (ground) atoms, and the non-circularity of paths according to the connection schemes in Figures 2(a)–(d) excludes mutual positive dependencies (between instances of `attack(X',Y',D)`). Hence, the modules induced by different instantiations of `board(n)` are compositional and can be joined to successively increase the board size.

The Python main routine in Lines 25–35 of Listing 11 controls the successive grounding and solving of a series of boards. To this end, an ordered list of non-overlapping integer intervals is to be provided on the command-line. For example, `-c calls="list((1,1),(3,5),(8,9))"` leads to successively solving the  $n$ -Queens problem for board sizes 1, 3, 4, 5, 8, and 9. As long as the upper limit of some interval is yet unreached, the board size is incremented by one in Line 30 and, in view of the `ground` instruction in Line 31, taken as a term for instantiating `board(n)`. However, solving is only invoked in Line 34 if the current size lies within the interval of interest. Provided that this is the case for any particular  $n \geq 1$ , the sequence of issued `ground` instructions makes sure that the current *clingo* state corresponds to the module obtained by instantiating and joining the subprograms `board(n/i)`, for  $1 \leq i \leq n$ , in increasing order. Since all ground rules accumulated in such a state are relevant (and not superseded by permanently falsifying the body) for  $n$ -Queens solving, there is no redundancy in instantiating `board(n/i)` for each  $1 \leq i \leq n$ , even when the provided integer intervals do not include  $i$  and  $i$ -Queens solving is skipped. For instance, `-c calls="list((1,1),(3,5),(8,9))"` specifies a series of six boards to solve, while the subprogram `board(n)` is successively instantiated with nine different terms for parameter  $n$ . In fact, the main routine in Lines 25–35 automates the assembly of subprograms needed to process an arbitrary yet increasing sequence of board sizes.

The result of running the program in Listing 11 with *clingo* is given in Listing 12.

```
$ clingo queens.lp -c calls="list((1,1),(3,5),(8,9))"
clingo version 4.5.4
Reading from queen.alt.lp
SIZE 1
```



Line 4	Line 12	Line 15	Line 17	Line 18	Line 20	Line 22
#external a(1,1,h).		{q(1,1)}.				:- not q(1,1), not a(1,1,h).
#external a(2,1,h).	t(2,1,1,1,h,2).	{q(2,1)}.	a(1,1,h) :- t(2,1,1,1,h,2), q(2,1).	a(1,1,h) :- t(2,1,1,1,h,2), a(2,1,h).	:- t(2,1,1,1,h,2), a(1,1,h), q(1,1).	
⋮	⋮	⋮	⋮	⋮	⋮	
#external a(n,1,h).	t(n,1,n-1,1,h,n).	{q(n,1)}.	a(n-1,1,h) :- t(n,1,n-1,1,h,n), q(n,1).	a(n-1,1,h) :- t(n,1,n-1,1,h,n), a(n,1,h).	:- t(n,1,n-1,1,h,n), a(n,1,h), q(n,1).	

Fig. 3. Incremental development of rules regarding the first row of a board.

```

Solving...
Answer: 1
queen(1,1)
SIZE 3
Solving...
SIZE 4
Solving...
Answer: 1
queen(2,1) queen(1,3) queen(4,2) queen(3,4)
SIZE 5
Solving...
Answer: 1
queen(2,1) queen(3,3) queen(1,4) queen(5,2) queen(4,5)
SIZE 8
Solving...
Answer: 1
queen(2,1) queen(1,4) queen(5,2) queen(3,5) queen(7,3) \
queen(6,7) queen(8,6) queen(4,8)
SIZE 9
Solving...
Answer: 1
queen(3,2) queen(1,3) queen(6,4) queen(5,6) queen(7,1) \
queen(2,7) queen(8,5) queen(4,8) queen(9,9)
SATISFIABLE

Models      : 5+
Calls       : 6
Time        : 0.013s (Solving: 0.00s 1st Model: 0.00s Unsat: 0.00s)
CPU Time    : 0.010s

```

Listing 12. Running the program in Listing 11 with *clingo*

### 5.3 Ricochet robots

In practice, ASP systems are embedded in encompassing software environments and thus need means for interaction. Multi-shot ASP solvers can address this by allowing a reactive procedure to loop on solving while acquiring changes to the problem specification. In this section, we want to illustrate this by modeling the popular board game of *Ricochet Robots*. Our particular focus lies on capturing the underlying round playing through the procedural-declarative interplay offered by *clingo*.

*Ricochet Robots* is a board game for multiple players designed by Alex Randolph.<sup>27</sup> A board consists of  $16 \times 16$  fields arranged in a grid structure having barriers between various neighboring fields (see Figures 4 and 5). Four differently colored robots roam across the board along either horizontally or vertically accessible fields, respectively. Each robot can thus move in four directions. A robot cannot stop its move until it either hits a barrier or another robot. The goal is to place a designated robot on a target location with a shortest sequence of moves. Often this involves

<sup>27</sup> [http://en.wikipedia.org/wiki/Ricochet\\_Robot](http://en.wikipedia.org/wiki/Ricochet_Robot)

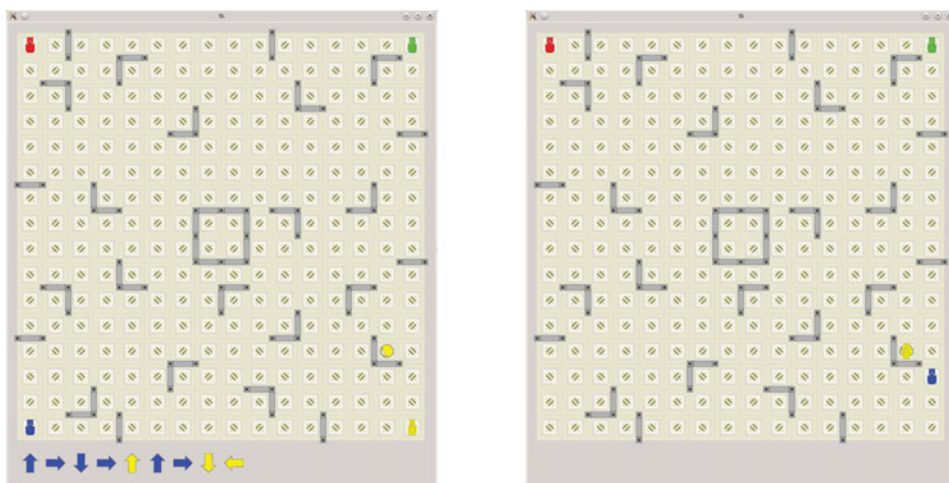


Fig. 4. Visualization of solving goal (13) from initially cornered robots.

moving several robots to establish temporary barriers. The game is played in rounds. At each round, a chip with a colored symbol indicating the target location is drawn. Then, the specific goal is to move the robot with the same color on this location. The player who reaches the goal with the fewest number of robot moves wins the chip. The next round is then played from the end configuration of the previous round. At the end, the player with most chips wins the game.

### 5.3.1 Encoding Ricochet Robots

The following encoding<sup>28</sup> and fact format follow the ones of Gebser *et al.* (2013).

An authentic board configuration of *Ricochet Robots* is shown in Figure 4 and represented as facts in Listing 13. The dimension of the board is fixed to 16 in Line 1. As put forward by Gebser *et al.* (2013), barriers are indicated by atoms with predicate `barrier/4`. The first two arguments give the field position, and the last two the orientation of the barrier, which is mostly east (1,0) or south (0,1).<sup>29</sup> For instance, the atom `barrier(2,1,1,0)` in Line 3 represents the vertical wall between the fields (2,1) and (3,1), and `barrier(5,1,0,1)` stands for the horizontal wall separating (5,1) from (5,2).

```
1 dim(1..16).
```

```
3 barrier( 2, 1, 1,0). barrier(13,11, 1,0). barrier( 9, 7,0, 1).
```

```
4 barrier(10, 1, 1,0). barrier(11,12, 1,0). barrier(11, 7,0, 1).
```

<sup>28</sup> Alternative ASP encodings of the game were studied by Gebser *et al.* (2013), and used for comparing various ASP solving techniques. More disparate encodings resulted from the ASP competition in 2013, where *Ricochet Robots* was included in the modeling track. ASP encodings and instances of *Ricochet Robots* are available at <https://potassco.org/doc/apps/2016/09/20/ricochet-robots.html>. There is also visualizer for the problem among the clingo examples: <https://github.com/potassco/clingo/tree/master/examples/clingo/robots>

<sup>29</sup> Symmetric barriers are handled by predicate `stop/4` in Lines 4 and 5 of Listing 15.

```

5 barrier( 4, 2, 1,0). barrier(14,13, 1,0). barrier(14, 7,0, 1).
6 barrier(14, 2, 1,0). barrier( 6,14, 1,0). barrier(16, 9,0, 1).
7 barrier( 2, 3, 1,0). barrier( 3,15, 1,0). barrier( 2,10,0, 1).
8 barrier(11, 3, 1,0). barrier(10,15, 1,0). barrier( 5,10,0, 1).
9 barrier( 7, 4, 1,0). barrier( 4,16, 1,0). barrier( 8,10,0,-1).
10 barrier( 3, 7, 1,0). barrier(12,16, 1,0). barrier( 9,10,0,-1).
11 barrier(14, 7, 1,0). barrier( 5, 1,0, 1). barrier( 9,10,0, 1).
12 barrier( 7, 8, 1,0). barrier(15, 1,0, 1). barrier(14,10,0, 1).
13 barrier(10, 8,-1,0). barrier( 2, 2,0, 1). barrier( 1,12,0, 1).
14 barrier(11, 8, 1,0). barrier(12, 3,0, 1). barrier(11,12,0, 1).
15 barrier( 7, 9, 1,0). barrier( 7, 4,0, 1). barrier( 7,13,0, 1).
16 barrier(10, 9,-1,0). barrier(16, 4,0, 1). barrier(15,13,0, 1).
17 barrier( 4,10, 1,0). barrier( 1, 6,0, 1). barrier(10,14,0, 1).
18 barrier( 2,11, 1,0). barrier( 4, 7,0, 1). barrier( 3,15,0, 1).
19 barrier( 8,11, 1,0). barrier( 8, 7,0, 1).

```

Listing 13. The Board (board.lp)

Listing 14 gives the 16 possible target locations printed on the game's carton board (cf. Lines 3 to 18). Each robot has four possible target locations, expressed by the ternary predicate `target`. Such a target is put in place via the unary predicate `goal` that associates a number with each location. The external declaration in Line 1 paves the way for fixing the target location from outside the solving process. For instance, setting `goal(13)` to true makes position (15, 13) a target location for the yellow robot.

```

1  #external goal(1..16).

3 target(red,    5, 2) :- goal(1).      % red moon
4 target(red,   15, 2) :- goal(2).      % red triangle
5 target(green,  2, 3) :- goal(3).      % green triangle
6 target(blue,  12, 3) :- goal(4).      % blue star
7 target(yellow, 7, 4) :- goal(5).      % yellow star
8 target(blue,   4, 7) :- goal(6).      % blue saturn
9 target(green, 14, 7) :- goal(7).      % green moon
10 target(yellow,11, 8) :- goal(8).      % yellow saturn
11 target(yellow, 5,10) :- goal(9).      % yellow moon
12 target(green,  2,11) :- goal(10).     % green star
13 target(red,   14,11) :- goal(11).     % red star
14 target(green, 11,12) :- goal(12).     % green saturn
15 target(yellow,15,13) :- goal(13).     % yellow star
16 target(blue,   7,14) :- goal(14).     % blue star
17 target(red,    3,15) :- goal(15).     % red saturn
18 target(blue,  10,15) :- goal(16).     % blue moon

20 robot(red;green;blue;yellow).
21 #external pos((red;green;blue;yellow),1..16,1..16).

```

Listing 14. Robots and targets (targets.lp)

Similarly, the initial robot positions can be set externally, as declared in Line 21. That is, each robot can be put at 256 different locations. On the left-hand side of Figure 4, we cornered all robots by setting `pos(red,1,1)`, `pos(blue,1,16)`, `pos(green,16,1)`, and `pos(yellow,16,16)` to true.

Finally, the encoding in Listing 15 gives a non-incremental encoding with a fixed horizon, following the one by Gebser *et al.* (2013, Listing 2).

```

1  time(1..horizon).
2  dir(-1,0;1,0;0,-1;0,1).

4  stop( DX, DY,X,   Y   ) :- barrier(X,Y,DX,DY).
5  stop(-DX,-DY,X+DX,Y+DY) :- stop(DX,DY,X,Y).

7  pos(R,X,Y,0) :- pos(R,X,Y).

9  1 { move(R,DX,DY,T) : robot(R), dir(DX,DY) } 1 :- time(T).
10 move(R,T) :- move(R,_,_,T).

12 halt(DX,DY,X-DX,Y-DY,T) :- pos(_,X,Y,T), dir(DX,DY),
13     dim(X-DX), dim(Y-DY), not stop(-DX,-DY,X,Y), T < horizon.

15 goto(R,DX,DY,X,Y,T) :- pos(R,X,Y,T), dir(DX,DY), T < horizon.
16 goto(R,DX,DY,X+DX,Y+DY,T) :- goto(R,DX,DY,X,Y,T),
17     dim(X+DX), dim(Y+DY), not stop(DX,DY,X,Y), not halt(DX,DY,X,Y,T).

19 pos(R,X,Y,T) :- move(R,DX,DY,T), goto(R,DX,DY,X,Y,T-1),
20     not goto(R,DX,DY,X+DX,Y+DY,T-1).
21 pos(R,X,Y,T) :- pos(R,X,Y,T-1), time(T), not move(R,T).

23 :- target(R,X,Y), not pos(R,X,Y,horizon).

25 #show move/4.

```

Listing 15. Simple encoding for *Ricochet Robots* (`ricochet.lp`)

The first lines in Listing 15 furnish domain definitions, fixing the sequence of time steps (`time/1`)<sup>30</sup> and two-dimensional representations of the four possible directions (`dir/2`). The constant `horizon` is expected to be provided via *clingo* option `-c` (e.g. “`-c horizon=20`”). Predicate `stop/4` is the symmetric version of `barrier/4` from above and identifies all blocked field transitions. The initial robot positions are fixed in Line 7 (in view of external input).

At each time step, some robot is moved in a direction (cf. Line 9). Such a move can be regarded as the composition of successive field transitions, captured by predicate `goto/6` (in Lines 15–17). To this end, predicate `halt/5` provides temporary barriers due to robots’ positions before the move. To be more precise, a robot moving in direction  $(DX, DY)$  must halt at field  $(X-DX, Y-DY)$  when some (other) robot is located at  $(X, Y)$ , and an instance of `halt(DX, DY, X-DX, Y-DY, T)` may provide information relevant to the move at step  $T+1$  if there is no barrier between  $(X-DX, Y-DY)$  and  $(X, Y)$ . Given this, the definition of `goto/6` starts at a robot’s position (in Line 15) and continues in direction  $(DX, DY)$  (in Lines 16–17) unless a barrier, a robot, or the board’s border is encountered. As this definition tolerates board traversals of length zero, `goto/6` is guaranteed to yield a successor position for any move of a robot  $R$  in direction  $(DX, DY)$ , so that the rule in

<sup>30</sup> The initial time point 0 is handled explicitly.

Lines 19–20 captures the effect of `move(R,DX,DY,T)`. Moreover, the frame axiom in Line 21 preserves the positions of unmoved robots, relying on the projection `move/2` (cf. Line 10).

Finally, we stipulate in Line 23 that a robot `R` must be at its target position `(X,Y)` at the last time point `horizon`. Adding directive “`#show move/4.`” further allows for projecting stable models onto the extension of the `move/4` predicate.

The encoding in Listing 15 allows us to decide whether a plan of length `horizon` exists. For computing a shortest plan, we may augment our decision encoding with an optimization directive. This can be accomplished by adding the part in Listing 16.

```
27 goon(T) :- target(R,X,Y), T = 0..horizon, not pos(R,X,Y,T).

29 :- move(R,DX,DY,T-1), time(T), not goon(T-1), not move(R,DX,DY,T).

31 #minimize{ 1,T : goon(T) }.
```

Listing 16. Encoding part for optimization (`optimization.lp`)

The rule in Line 27 indicates whether some goal condition is (not) established at a time point. Once the goal is established, the additional integrity constraint in Line 29 ensures that it remains satisfied by enforcing that the goal-achieving move is repeated at later steps (without altering robots’ positions). Note that the `#minimize` directive in Line 31 aims at few instances of `goon/1`, corresponding to an early establishment of the goal, while further repetitions of the goal-achieving move are ignored. Our extended encoding allows for computing a shortest plan of length bounded by `horizon`. If there is no such plan, the problem can be posed again with an enlarged `horizon`. For computing a shortest plan in an unbounded fashion, we can take advantage of incremental ASP solving, as illustrated in Section 5.1.

Apart from the two external directives that allow us to vary initial robot and target positions, the four programs constitute an ordinary ASP formalization of a *Ricochet Robots* instance. To illustrate this, let us override the external directives by adding facts accounting for the robot and target positions on the left hand side of Figure 4. The corresponding call of *clingo* is shown in Listing 17.<sup>31</sup>

```
1 $ clingo board.lp targets.lp ricochet.lp optimization.lp \
2     -c horizon=10 \
3     <(echo "pos(red,1,1).    pos(green,16,1).    \
4         pos(blue,1,16). pos(yellow,16,16).    \
5         goal(13).")
```

Listing 17. One-shot solving with *clingo*

```
1 move(blue,0,-1,1)    move(blue,1,0,2)    move(blue,0,1,3)    \
2 move(blue,1,0,4)    move(yellow,0,-1,5)  move(blue,0,-1,6)  \
3 move(blue,1,0,7)    move(yellow,0,1,8)  move(yellow,-1,0,9) \
4 move(yellow,-1,0,10)
```

Listing 18. Stable model projected onto the extension of the `move/4` predicate

The resulting one-shot solving process yields a(n optimal) stable model containing the extension of the `move/4` predicate given in Listing 18. The `move` atoms in Lines 1–4

<sup>31</sup> Note that rather than using input redirection, we also could have passed the five facts via a file.

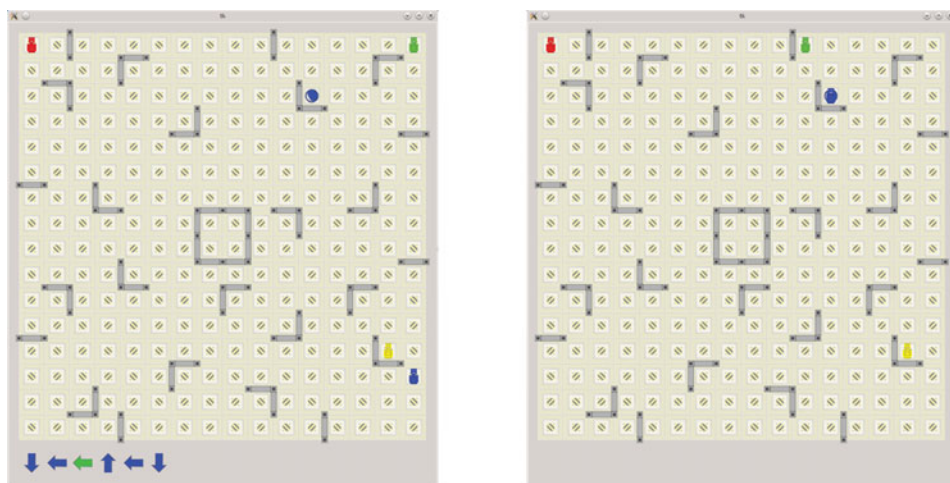


Fig. 5. Visualization of solving goal (4) from robot positions after having solved goal (13).

of Listing 18 correspond to the plan indicated by the colored arrows at the bottom of the left hand side of Figure 4. That is, the blue robot starts by going north, east, south, and east, then the yellow one goes north, the blue one resumes and goes north and east, before finally the yellow robot goes south (bouncing off the blue one) and lands on the target by going west. This leads to the situation depicted on the right hand side of Figure 4. Note that the tenth move (in Line 4) is redundant since it merely replicates the previous one because the goal was already reached after nine steps.

### 5.3.2 Playing in rounds

*Ricochet Robots* is played in rounds. Hence, the next goal must be reached with robots placed at the positions resulting from the previous round. For example, when pursuing goal (4) in the next round, the robots must start from the end positions given on the right hand side of Figure 4. The resulting configuration is shown on the left hand side of Figure 5. For one-shot solving, we would relaunch *clingo* from scratch as shown in Listing 17, yet by accounting for the new target and robot positions by replacing Lines 3–5 of Listing 17 by the following ones.

```

3      <(echo "pos (red, 1, 1) .   pos (green, 16, 1) .   \
4          pos (blue, 16, 10) . pos (yellow, 15, 13) . \
5          goal (4) . "

```

Unlike this, our multi-shot approach to playing in rounds relies upon a single<sup>32</sup> operational *clingo* control object that we use in a simple loop:

- (1) Create an operational control object (containing a grounder and a solver object)
- (2) Load and ground the programs in Listing 13, 14, 15, and optionally 16 (relative to some fixed horizon) within the control object

<sup>32</sup> In general, multiple such control objects can be created and made to interact via Python.

- (3) While there is a goal, do the following
  - (a) Enforce the initial robot positions
  - (b) Enforce the current goal
  - (c) Solve the logic program contained in the control object

The control loop is implemented in Python by means of *clingo*'s Python API. This module provides grounding and solving functionalities.<sup>33</sup> As mentioned in Section 2, both modules support (almost) literal counterparts to “Create,” “Load,” “Ground,” and “Solve.” The “enforcement” of robot and target positions is more complex, as it involves changing the truth values of externally controlled atoms (mimicking the insertion and deletion of atoms, respectively).

The resulting Python program is given in Listing 19. Line 1 imports the `clingo` module. We are only using three classes from the module, which we directly pull into the global namespace to avoid qualification with “`clingo.`” and so to keep the code compact.

Lines 3–34 show the `Player` class. This class encapsulates all state information including *clingo*'s `Control` object that in turn holds the state of the underlying grounder and solver. In the `Player`'s `__init__` function (similar to a constructor in other object-oriented languages) the following member variables are initialized:

**last\_positions** This variable is initialized upon construction with the starting positions of the robots. During the progression of the game, this variable holds the initial starting positions of the robots for each turn.

**last\_solution** This variable holds the last solution of a search call.

**undo\_external** We want to successively solve a sequence of goals. In each step, a goal has to be reached from different starting positions. This variable holds a list containing the current goal and starting positions that have to be cleared upon the next step.

**horizon** We are using a bounded encoding. This (Python) variable holds the maximum number of moves to find a solution for a given step.

**ctl** This variable holds the actual object providing an interface to the grounder and solver. It holds all state information necessary for multi-shot solving along with heuristic information gathered during solving.

As shown in Lines 4–13, the constructor takes the `horizon`, initial robot positions, and the files containing the various logic programs. *clingo*'s `Control` object is created in Lines 9–10 by passing the option `-c` to replace the logic program constant `horizon` by the value of the Python variable `horizon` during grounding. Finally, the constructor loads all files and grounds the entire logic program in Lines 11–13. Recall from Section 2 that all rules outside the scope of `#program` directives belong to the base program. Note also that this is the only time grounding happens because the encoding is bounded. All following solving steps are configured exclusively via manipulating external atoms.

<sup>33</sup> An analogous module is available for Lua.



```

1  from clingo import Control, Model, Function
2
3  class Player:
4      def __init__(self, horizon, positions, files):
5          self.last_positions = positions
6          self.last_solution = None
7          self.undo_external = []
8          self.horizon = horizon
9          self.ctl = Control(
10             ['-c', 'horizon={0}'.format(self.horizon)])
11         for x in files:
12             self.ctl.load(x)
13         self.ctl.ground(["base", []])
14
15     def solve(self, goal):
16         for x in self.undo_external:
17             self.ctl.assign_external(x, False)
18         self.undo_external = []
19         for x in self.last_positions + [goal]:
20             self.ctl.assign_external(x, True)
21             self.undo_external.append(x)
22         self.last_solution = None
23         self.ctl.solve(on_model=self.on_model)
24         return self.last_solution
25
26     def on_model(self, model):
27         self.last_solution = model.symbols(atoms=True)
28         self.last_positions = []
29         for atom in model.symbols(atoms=True):
30             if (atom.name == "pos" and
31                 len(atom.arguments) == 4 and
32                 atom.arguments[3].number == self.horizon):
33                 self.last_positions.append(
34                     Function("pos", atom.arguments[:-1]))
35
36 horizon = 15
37 encodings = ["board.lp", "targets.lp", "ricochet.lp", "optimization.lp"]
38 positions = [Function("pos", [Function("red"), 1, 1]),
39              Function("pos", [Function("blue"), 1, 16]),
40              Function("pos", [Function("green"), 16, 1]),
41              Function("pos", [Function("yellow"), 16, 16])]
42 sequence = [Function("goal", [13]),
43            Function("goal", [4]),
44            Function("goal", [7])]
45
46 player = Player(horizon, positions, encodings)
47 for goal in sequence:
48     print player.solve(goal)

```

Listing 19. The Ricochet Robot Player (ricochet.py)

The `solve` method in Lines 15–24 starts with initializing the search for the solution to the new goal. To this end, it first undoes in Lines 16–17 the previous goal and starting positions stored in `undo_external` by assigning `False` to the respective atoms. In the following Lines 19–21, the next step is initialized by assigning `True` to the current goal along with the last robot positions; these are also stored in `undo_external` so that they can be taken back afterwards. Finally, the `solve` method calls *clingo*'s `ctl.solve` to initiate the search. The result is captured in variable `last_solution`. Note that the call to `ctl.solve` takes `ctl.on_model` as (keyword) argument, which is called whenever a model is found. In other words, `on_model` acts as a callback for intercepting models. Finally, variable `last_solution` is returned at the end of the method.

The last function of the `Player` class is the `on_model` callback. As mentioned, it intercepts the (final) models computed by the solver, which can then be inspected via the functions of the `Model` class. At first, it stores the shown atoms in variable `last_solution` in Line 27.<sup>34</sup> The remainder of the `on_model` callback extracts the final robot positions from the stable model. For that, it loops in Lines 29–34 over the full set of atoms in the model and checks whether their signatures match. That is, if an atom is formed from predicate `pos/4`, and its fourth argument equals the horizon, then it is appended to the list of `last_positions` after stripping its time step from its arguments.

As an example, consider `pos(yellow, 15, 13, 20)`, say the final position of the yellow robot on the right hand side of Figure 4 at an horizon of 20. This leads to the addition of `pos(yellow, 15, 13)` to the `last_positions`. Note that `pos(yellow, 15, 13)` is declared an external atom in Line 21 of Listing 14. For playing the next round, we can thus make it `True` in Line 20 of Listing 19. And when solving, the rule in Line 7 of Listing 15 allows us to derive `pos(yellow, 15, 13, 0)` and makes it the new starting position of the yellow robot, as shown on the left hand side of Figure 5.

Lines 36–44 show the code for configuring the player. They set the search horizon, the encodings to solve with, and the initial positions in form of *clingo* terms. Furthermore, we fix a sequence of goals in Lines 42–44. In a more realistic setting, either some user interaction or a random sequence might be generated to emulate arbitrary draws.

```

1 $ python ricochet.py
2 [move(red,0,1,1), move(red,1,0,2), move(red,0,1,3), ...]
3 [move(blue,0,-1,1), move(blue,1,0,2), move(blue,0,1,3), ...]
4 [move(green,0,1,1), move(green,1,0,2), move(green,1,0,3), ...]

```

Listing 20. Multi-shot solving with *clingo*'s Python API

Finally, Lines 46–48 implement the search for sequences of moves that solve the configuration given above. For each goal in the sequence, a solution is plainly printed, as engaged in Line 48. The three lists in Listing 20 represent solutions to the three goals in Lines 42–44. The *clingo* library does not foresee any output,

<sup>34</sup> In view of “`#show move/4.`” in Listing 15, this only involves instances of `move/4`, while all true atoms are included via the argument `Model.ATOMS` in Line 29.

which must thus be handled by the scripting language. Note also that the first list represents an alternative solution to the one given in Listing 18.

### 5.4 Optimization

Another innovative feature of *clingo* is its incremental optimization. This allows for adapting objective functions along the evolution of a program at hand. A simple example is the search for shortest plans when increasing the horizon in non-consecutive steps. To see this, recall that literals in minimize statements (and analogously weak constraints) are supplied with a sequence of terms of the form  $w@p,\vec{t}$ , where  $w$  and  $p$  are integers providing a weight and a priority level and  $\vec{t}$  is a sequence of terms (cf. Calimeri *et al.* (2012)). As an example, consider the subprogram<sup>35</sup>:

```
#program cumulativeObjective(t) .
#minimize{ W@P,X,Y,t : move(X,Y,W,P,t) } .
```

When grounding and solving `cumulativeObjective(t)` for successive values of  $t$ , the solver's objective function (per priority level  $P$ ) is gradually extended with new atoms over `move/5`, and all previous ones are kept.

For enabling the removal of literals from objective functions, we can use externals:

```
#program volatileObjective(t) .
#external activateObjective(t) .
#minimize{ W@P,X,Y,t : move(X,Y,W,P,t), activateObjective(t) } .
```

The subprogram `volatileObjective(t)` behaves like `cumulativeObjective(t)` as long as the external atom `activateObjective(t)` is true. Once it is set to false, all atoms over `move/5` with the corresponding term for  $t$  are dismissed from objective functions.

## 6 Application program interfaces

This section provides some further selected functionalities of *clingo*'s APIs; detailed descriptions can be found at [potassco.org](http://potassco.org). Currently, *clingo* provides APIs in C, C++, Lua, and Python, all sharing the same functionality. A tutorial on using the Python API for multi-shot and theory solving was given by Kaminski *et al.* (2017).

The theory reasoning capabilities of *clingo* are described by Gebser *et al.* (2016). In brief, *clingo* provides generic means for incorporating theory reasoning. They span from theory grammars for seamlessly extending its input language with theory expressions to a simple interface for integrating theory propagators into its solver component. Multi-shot solving for selected theories is described by Banbara *et al.* (2017) and Janhunen *et al.* (2017).

The central role in multi-shot solving is played by control objects capturing the system states of grounders and solvers, as introduced in Section 4.5. While the control object created by invoking *clingo* from the command line is passed as argument to

<sup>35</sup> The same applies to a weak constraint of form “:  $\sim\text{move}(X,Y,W,P,t)$  . [W@P,X,Y,t].”

the main routine, further such objects can be created with the constructor `Control`. Examples for both settings can be found in Line 8 of Listing 7 and Lines 9 and 10 of Listing 19, respectively. In general, this allows multiple independent *clingo* objects to coexist and communicate with each other.

*clingo*'s interface can be structured into three parts.

*Parsing.* A simple but very useful feature of *clingo*'s API is that it allows us to leverage the parser of its grounding component *gringo* to obtain an abstract syntax tree (AST) of the non-ground program. More precisely, the interface allows for both obtaining an AST of a full-fledged logic program and adding such a program in form of an AST. This provides an easy way of applying program transformations on the non-ground level without any burden of parsing input programs in their full generality. This feature is exploited by the systems *asprin* (Brewka et al. 2015b) for expressing preferences and *anthem*<sup>36</sup> for formula extraction.

*Grounding.* A basic functionality of *clingo* objects is to incrementally augment non-ground programs by loading programs from file or adding them in string form. An example of the former can be found in Line 12 of Listing 19; the latter is accomplished by the counterpart of *add*, defined in Section 4.5. Notably, this functionality allows for adding dynamically generated programs.

More fine-grained control is provided by the low level part of the interface. This allows, for instance, for inspecting the result of grounding and adding ground rules in the intermediate ASP format *aspif* (cf. Kaminski et al. (2017)). In this way, one can iterate over all ground atoms and inspect them individually, or implement eager constraint translations by adding new (theory) atoms on demand. Also, symbols can be injected during grounding via external functions, similar to value invention (Calimeri et al. 2007).

*Solving.* The principal `solve` method can be shaped in various ways. For instance, we have seen in Listing 19 how `on_model` can act as a callback for intercepting models. More precisely, for each stable model found during a call to `solve(on_model=f)`, a model object is passed to function `f`, whose implementation can then access and inspect the model. An example consists of the addition of constraints whenever a model is found,<sup>37</sup> as in optimization tasks, or final tests on model candidates. Similarly, `solve` can be supplied with assumptions, as detailed in (19). For instance, the call `solve(assumptions=[(Function("a"), True)])` only admits stable models containing atom `a`. Moreover, *clingo* provides an asynchronous interface, which is particularly useful in reactive settings. Here, solving is done in the background and interruptible at any time. For example, this allows to accommodate scenarios where agents have to stay responsive even though solving has not yet finished. Finally, it is worth mentioning that dedicated parts of the API allow for configuring search and extracting solver statistics. In combination with the aforementioned `on_model`

<sup>36</sup> <https://github.com/potassco/anthem>

<sup>37</sup> A more sophisticated way is to use theory propagators adding constraints not just when a model is found but also during the solver's propagation; see Gebser et al. (2016) for details.

callback this allows for re-configuring search in view of the statistics gathered during the search for the last model.

## 7 Experiments

The computational advantage of multi-shot solving lies in its avoidance of redundancies otherwise caused by relaunching grounder and solver programs. Since the substantial savings on grounding intense benchmarks have already been demonstrated (Gebser *et al.* 2008), we focus our empirical analysis on the impact of our approach on solving. In particular, we want to investigate in how far multi-shot solving can benefit from the learning capacities of modern ASP solvers and the reuse of already gathered heuristic scores. To this end, we empirically evaluate the impact of *clingo*'s multi-shot solving capacities on three planning benchmarks:<sup>38</sup> Towers of Hanoi (cf. Section 5.1), Ricochet Robots (cf. Section 5.3), and ASP encodings of PDDL problems (Dimopoulos *et al.* 2017).<sup>39</sup> For either benchmark, we let *clingo* version 4.5.4 search for a shortest plan by incrementally extending the horizon until the first plan is found. In particular, we consider multi-shot solving by means of *clingo*'s built-in incremental mode (invoked by “`#include <incmode>.`”) in four different settings:

- *Multi*: keeping recorded nogoods as well as heuristic values between solver calls.
- *Multi -heuristic*: keeping recorded nogoods, but not heuristic values, between solver calls.
- *Multi -nogoods*: keeping heuristic values, but not recorded nogoods, between solver calls.
- *Multi-heuristic-nogoods*: keeping neither recorded nogoods nor heuristic values between solver calls.

We contrast these four settings to the traditional single-shot approach, denoted by *single*, where *clingo* performs grounding and solving from scratch for each planning horizon. The experiments were run sequentially on a Linux machine equipped with Xeon E5520 2.27 GHz processors, limiting wall-clock time to 3,000 seconds per run without imposing any (effective) memory limit. *clingo* was run in all experiments in its default configuration except for the option `--forget-on-step` that allows for configuring the four settings above.

We portray our results in terms of cactus plots, in which the *x*-axes list instances ordered by time (and conflicts) and the *y*-axes reflect times and conflicts, respectively. The magnitude of the latter are given on top of the *y*-axis.

<sup>38</sup> The benchmarks are available at <http://www.cs.uni-potsdam.de/wv/clingo/benchmark-2017-05-26.tar.xz>

<sup>39</sup> PDDL stands for Planning Domain Definition Language and should indicate that our benchmarks were obtained by translating benchmarks originally specified in PDDL and used by the planning community.

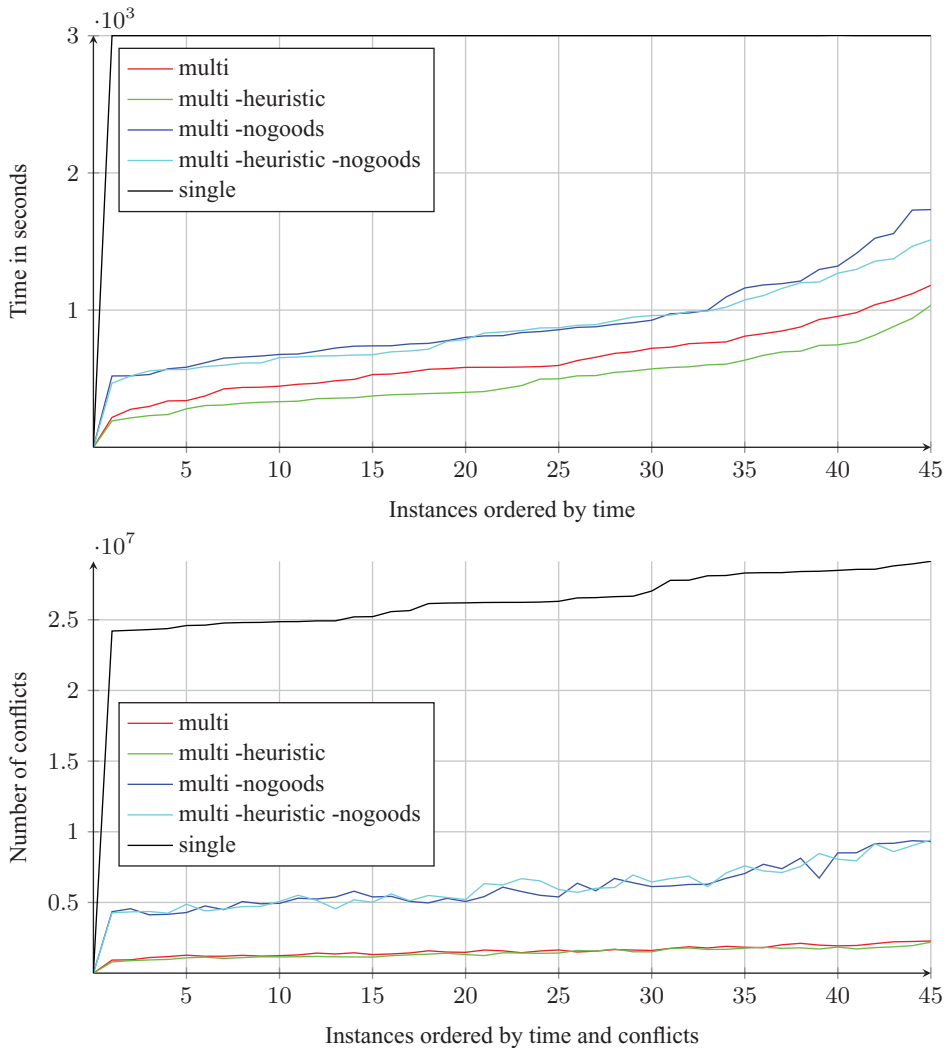


Fig. 6. Cactus plots for Towers of Hanoi benchmark.

### 7.1 Towers of Hanoi

The upper plot in Figure 6 displays runtimes for each *clingo* setting in increasing order over 45 instances of the Towers of Hanoi benchmark. Most apparently, we observe that single-shot solving performed in the *single* setting fails to complete any of the instances within the allotted 3,000 seconds. This clearly shows that the redundancy of relaunching grounding and solving processes from scratch for each horizon incurs non-negligible overhead here.

Somewhat unexpectedly, the multi-shot solving approaches in which recorded nogoods are discarded between successive solver calls, viz. *multi-nogoods* and *multi-heuristic-nogoods*, perform much better than single-shot solving and complete all 45 instances within the time limit. On the one hand, this advantage is owed to

incremental grounding, adding only new rule instances when switching from one horizon to the next. On the other hand, we verified that solving steps take the major share of runtime with each setting, and the numbers of conflicts plotted in the lower part of Figure 6 exhibit substantial search reductions in multi-shot solving, even when neither recorded nogoods nor heuristic values are kept. In fact, at the implementation level *clingo* asserts unary nogoods and stores binary as well as ternary nogoods persistently in dedicated data structures. Hence, such short nogoods remain available in multi-shot solving regardless of settings and explain the gap to iterated single-shot solving, where respective information has to be repeatedly retrieved from search conflicts at each solving step.

Keeping also long nogoods between solver calls, as done in the *multi* and *multi-heuristic* settings, further reduces search conflicts and runtime by a factor of about 5 or 1.5, respectively. This indicates a trade-off between memory demands and search savings due to recorded nogoods, where the savings outweigh the overhead on the Towers of Hanoi benchmark. Unlike that, keeping heuristic values does not pay off here, and the two *-heuristic* settings save some fraction of runtime in comparison to their counterparts passing such values between solver calls. Although respective gaps are modest, we conclude that biasing search to proceed as in previous solving steps does not provide shortcuts for problems with an extended planning horizon.

## 7.2 Ricochet robots

The behavior of single- and multi-shot solving approaches on 38 instances of the Ricochet Robots benchmark, plotted in Figure 7, parallels the previous observations. While the *single* setting is able to complete 22 of the instances within the given time, the most successful multi-shot solving approach, *multi-heuristic-nogoods*, solves all of them. In contrast to Towers of Hanoi above, keeping long nogoods does not pay off here, as the lower plot in Figure 7 shows that they do not significantly reduce the search conflicts at solving steps with an extended horizon. Hence, the two *-nogoods* settings are ahead in terms of runtime as well as solved instances displayed in the upper part of Figure 7. Moreover, we see that keeping or discarding heuristic values, the latter denoted by *-heuristic*, does not make much difference, which again exposes that biasing the search process in view of previous solving steps does not necessarily help for problem extensions.

## 7.3 PDDL problems

The performance results displayed in Figure 8, summarizing 20 instances obtained by translating planning problems from PDDL, also exhibit a significant gap separating single-shot from multi-shot solving in its four settings. The behavior of the latter differs primarily w.r.t. the treatment of long nogoods, where only the two *-nogoods* settings that discard them between solver calls complete all instances in time. In fact, comparing runtimes in the upper and numbers of conflicts in the lower part of Figure 8, it turns out that keeping such long nogoods incurs overhead without bringing about (consistent) search savings in return. The reduced numbers of conflicts

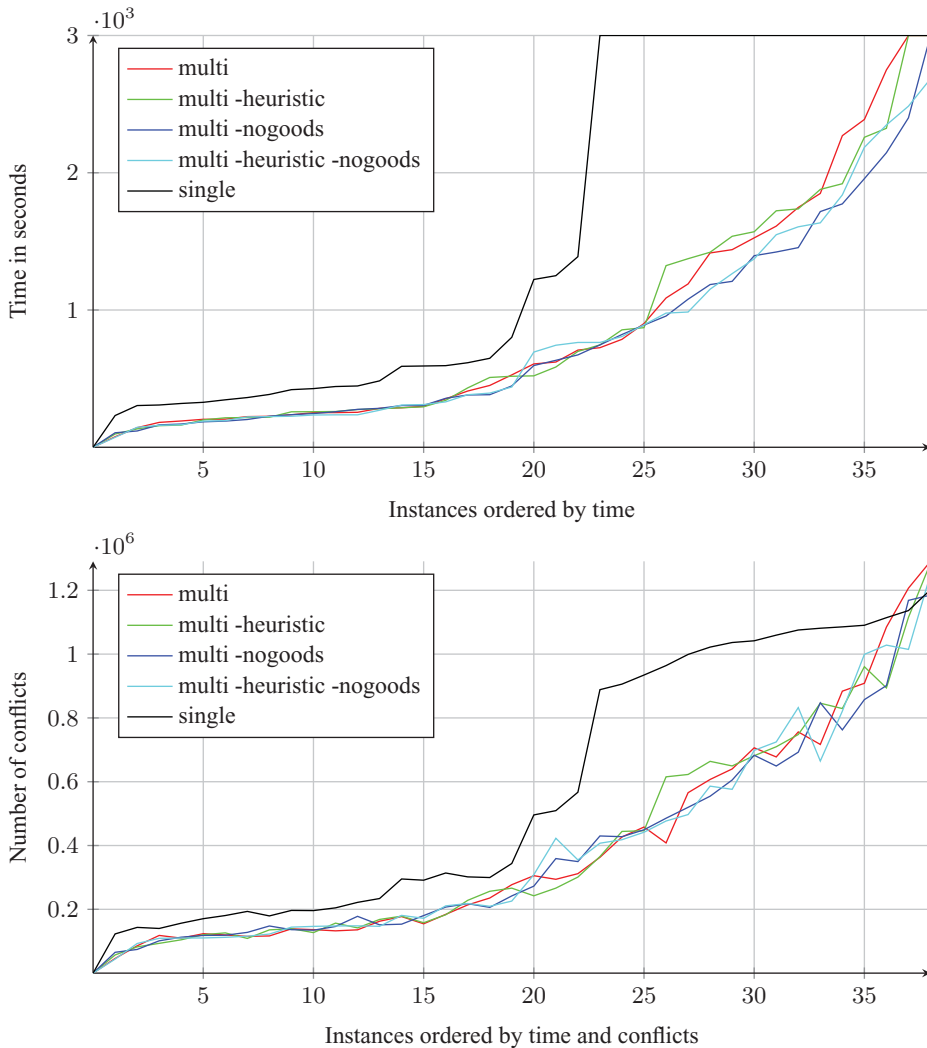


Fig. 7. Cactus plots for Ricochet Robots benchmark.

relative to single-shot solving nevertheless indicate a substantial added value of keeping some nogoods, in particular, short ones, between solver calls.

#### 7.4 Observations

On all benchmarks the number of conflicts is lower in the *multi* settings compared to the *single* one. This is due to the reuse of nogoods learnt from previous solving steps. In fact, the treatment of nogoods had the greatest influence on the runtime in the *multi* settings. As regards long nogoods, the sometimes unproportional overhead of keeping them, observed on the Ricochet Robots benchmark and PDDL problems, suggests to strive in the future for adaptive filtering mechanisms (beyond regular nogood deletion) assessing the relevance of recorded nogoods from previous solving



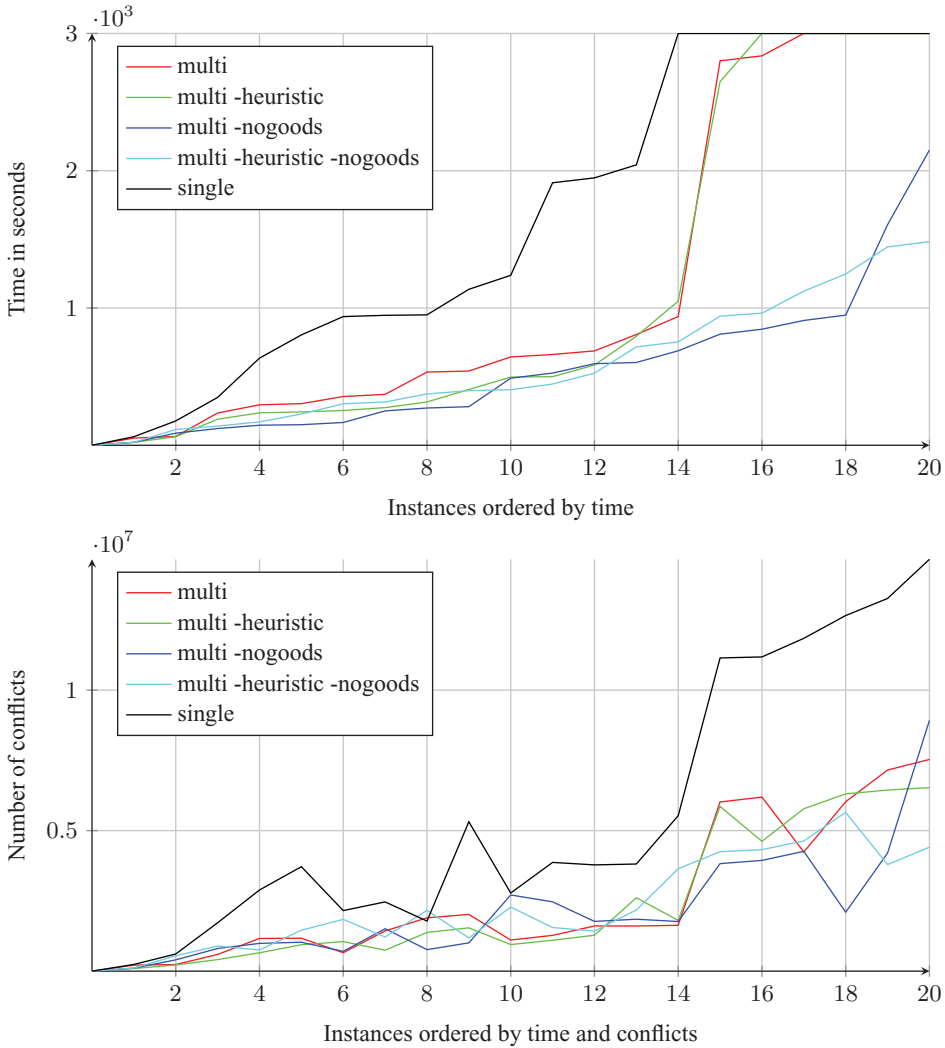


Fig. 8. Cactus plots for PDDL benchmark.

steps. Unlike that, *clingo* settings that differ in the treatment of heuristic values only, i.e., pairs denoted with or without *-heuristic*, exhibit comparable behavior on all three investigated benchmark problems, thus indicating that recorded nogoods play a more important role upon successive solver calls.

Our benchmarks were not very memory demanding. The peak memory consumption over all benchmarks was 187 MB when multi-shot solving and only 150 MB in the single-shot setting. Even though it might seem that the single-shot approach has a smaller memory footprint, the multi-shot settings probably just held larger databases of learnt clauses leading to a slightly higher memory consumption. With the selected benchmarks, the memory used by the grounding component is negligible.

## 8 Related work

Although *clingo* (Gebser et al. 2011) already featured Lua as an embedded scripting language up to series 3, its usage was limited to (deterministic) computations during grounding; neither were library functions furnished by *clingo* 3.

Of particular interest is *dlvhex* (Eiter et al. 2012), an ASP system aiming at the integration of external computation sources. For this purpose, *dlvhex* relies on higher-order logic programs using external higher-order atoms for software interoperability. Such external atoms should not be confused with *clingo*'s `#external` directive because they are evaluated via procedural means during solving. Given this, *dlvhex* can be seen as an *ASP modulo Theory* solver, similar to SAT modulo Theory solvers (Nieuwenhuis et al. 2006). In fact, *dlvhex* is build upon *clingo* and follows the design of the *ASP modulo CSP* solver *clingcon* (Ostrowski and Schaub 2012) in communicating with external “oracles” through *clasp*'s post propagation mechanism. In this way, theory solvers are tightly integrated into the ASP system and have access to the solver's partial assignments. Unlike this, multi-shot solving only provides access to total (stable) assignments. This is why *clingo* also offers full-fledged theory reasoning capabilities, dealing with partial assignments (Gebser et al. 2016; Kaminski et al. 2017). Clearly, the above considerations also apply to extensions of *dlvhex*, such as *acthex* (Fink et al. 2013). Furthermore, *jdlv* (Febbraro et al. 2012) encapsulates the *dlv* system to facilitate one-shot ASP solving in Java environments by providing means to generate and process logic programs, and to afterwards extract their stable models. *embasp* (Fuscà et al. 2016) provides a more recent and more general environment for embedding ASP systems, including *clingo* and *dlv*, into external systems. Meanwhile the ASP solver *wasp* (Alviano et al. 2015) also features a foreign language API, yet restricted to solving functionalities. More precisely, it provides low-level functionalities to customize heuristics and propagation (Dodaro et al. 2016a; Dodaro et al. 2016b).

The procedural attachment to the *idp* system (De Pooter et al. 2013; De Cat et al. 2014) builds on interfaces to C++ and Lua. Like *clingo*, it allows for evaluating functions during grounding, calling the grounder and solver multiple times, inspecting solutions, and reacting to external input after search. The emphasis, however, lies on high-level control blending in with *idp*'s modeling language, while *clingo* offers more fine-grained control over the grounding and solving process, particularly aiming at a flexible incremental assembly of programs from subprograms.

In SAT, incremental solver interfaces from low-level APIs are common practice. Pioneering work was done in *minisat* (Eén and Sörensson 2004), furnishing a C++ interface for solving under assumptions. In fact, the *clasp* library underlying *clingo* builds upon this functionality to implement incremental search (see Gebser et al. (2008)). Given that SAT deals with propositional formulas only, solvers and their APIs lack support for modeling languages and grounding. Unlike this, the SAT modulo Theory solver *z3* (de Moura and Bjørner 2008) comes with a Python API that, similar to *clingo*, provides a library for controlling the solver as well as language bindings for constraint handling. In this way, Python can be used as a modeling language for *z3*.

## 9 Conclusion

The *clingo* system complements ASP's declarative input language by control capacities expressed either by embedded scripting languages or by importing *clingo* modules into imperative programs. This is accomplished within a single integrated ASP grounding and solving process in which a logic program may evolve over time. The addition, deletion, and replacement of programs are controlled procedurally by means of *clingo*'s API. Applications that cannot be captured with the standard one-shot approach of ASP but that require evolving logic programs are manifold. Examples include unrolling a transition function as in planning, interacting with an environment as in assisted living, robotics, or stream reasoning, interacting with a user exploring a domain, theory solving, and advanced forms of search. Addressing these demands by providing a high-level API yields a generic and transparent approach. Unlike this, previous systems, like *iclingo* and *oclingo*, had a dedicated purpose involving rigid control procedures buried in monolithic programs. Rather than that, the basic technology of *clingo* allows us to instantiate subprograms in-between solver invocations in a fully customizable way. On the declarative side, the availability of program parameters and the embedding of `#external` directives into the grounding process provide us with a great flexibility in modeling schematic subprograms. In addition, the possibility of assigning input atoms facilitates the implementation of applications such as query answering (Gebser *et al.* 2013) or sliding window reasoning (Gebser *et al.* 2012), as truth values can now be switched without modifying logic programs.

The semantic underpinnings of our framework in terms of module theory capture the dynamic combination of logic programs in a generic way. Although this eases the modular composition of data structures, other choices are possible at the cost of higher maintenance. Note that the difficulty of composing subprograms in ASP is due to its nonmonotonic nature; this is much easier in monotonic approach such as SAT. Finally, it is interesting future work to investigate how dedicated change operations that were so far only of theoretic interest, like updating (Alferes *et al.* 2002), forgetting (Zhang and Foo 2006), revising (Delgrande *et al.* 2008), or merging (Delgrande *et al.* 2009) etc., can be put into practice within this framework. A first attempt at capturing the update of logic programs was made by Sabuncu and Leite (2017).

The input language of *clingo* extends the *ASP-Core-2* standard (Calimeri *et al.* 2012) and has meanwhile been put in its entirety on solid semantic foundations (Gebser *et al.* 2015). Although we have presented *clingo* for normal logic programs, we mention that it accepts (extended) disjunctive logic programs processed via the multi-threaded solving approach of *clasp* (Gebser *et al.* 2013). Since *clingo* embeds *clasp* series 3, it moreover features domain-specific heuristics (Gebser *et al.* 2013) and optimization using unsatisfiable cores (Andres *et al.* 2012). *clingo* is freely available at [potassco.org](http://potassco.org), and its releases include many best practice examples illustrating the aforementioned application scenarios.

Since its first release and accompanying publication (Gebser *et al.* 2014), *clingo*'s multi-shot solving framework has been used for implementing several ASP-based

reasoning systems, such as *asprin* (Brewka et al. 2015a; Brewka et al. 2015b), *aspic* (Gebser et al. 2013), *rosoclingo* (Andres et al. 2015), and *dflat* (Abseher et al. 2014); various forms of aggregates were implemented with it by Alviano et al. (2015) and Alviano and Leone (2015). As well, *dlvhex* (Eiter et al. 2012) builds upon *clingo* and its versatile API. This already hints at the potential impact of *clingo*'s multi-shot ASP solving framework, and we believe that it constitutes a step towards putting more and more applications into the reach of ASP.

Multi-shot ASP solving broadens the spectrum of applications of ASP. This also brings about the new user profile of *ASP engineering* that combines ASP modeling with traditional programming for controlling an ASP solving process. This may lead to generic advanced forms of ASP solving such as the incremental approach in Section 5.1 or be restricted to customized settings as with the Ricochet Robots game in Section 5.3. Multi-shot solving enables users to engineer such novel declarative systems on top of ASP. We believe that this new engineering facet is crucial to putting ASP into practice.

## References

- ABITEBOUL, S., HULL, R. AND VIANU, V. 1995. *Foundations of Databases*. Addison-Wesley.
- ABSEHER, M., BLIEM, B., CHARWAT, G., DUSBERGER, F., HECHER, M. AND WOLTRAN, S. 2014. The D-FLAT system for dynamic programming on tree decompositions. In *Proc. of the 14th European Conference on Logics in Artificial Intelligence (JELIA'14)*, E. Fermé and J. Leite, Eds. Lecture Notes in Artificial Intelligence, vol. 8761. Springer-Verlag, 558–572.
- ALFERES, J., PEREIRA, L., PRZYMUSINSKA, H. AND PRZYMUSINSKI, T. 2002. LUPS: A language for updating logic programs. *Artificial Intelligence* 138, 1–2, 87–116.
- ALVIANO, M., DODARO, C., LEONE, N. AND RICCA, F. 2015. Advances in WASP. *Proceedings of the 13th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'15)*, Calimeri, F., Ianni, G. and Truszczyński, M., Eds. Lecture Notes in Artificial Intelligence, vol. 9345. Springer-Verlag, 40–54.
- ALVIANO, M., FABER, W. AND GEBSER, M. 2015. Rewriting recursive aggregates in answer set programming: Back to monotonicity. *Theory and Practice of Logic Programming* 15, 4–5, 559–573. URL: <http://arxiv.org/abs/1507.03923>.
- 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.
- ANDRES, B., KAUFMANN, B., MATHEIS, O. AND SCHAUB, T. 2012. Unsatisfiability-based optimization in clasp. In *Proc. of Technical Communications of the 28th International Conference on Logic Programming (ICLP'12)*, A. Dovier and V. Santos Costa, Eds., Leibniz International Proceedings in Informatics (LIPIcs), vol. 17. University of Potsdam, 212–221.
- ANDRES, B., RAJARATNAM, D., SABUNCU, O. AND SCHAUB, T. 2015. Integrating ASP into ROS for reasoning in robots. See *Proceedings of the 13th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'15)*, Calimeri, F., Ianni, G. and Truszczyński, M., Eds. Lecture Notes in Artificial Intelligence, vol. 9345. Springer-Verlag, 69–82.
- BALDUCCINI, M. AND JANHUNEN, T., Eds. 2017. *Proceedings of the 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'17)*. Lecture Notes in Artificial Intelligence, vol. 10377. Springer-Verlag.

- BANBARA, M., KAUFMANN, B., OSTROWSKI, M. AND SCHAUB, T. 2017. Clingcon: The next generation. *Theory and Practice of Logic Programming* 17, 4, 408–461.
- BARAL, C. 2003. *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge University Press.
- BREWKA, G., DELGRANDE, J., ROMERO, J. AND SCHAUB, T. 2015a. asprin: Customizing answer set preferences without a headache. In *Proc. of the 29th National Conference on Artificial Intelligence (AAAI'15)*, B. Bonet and S. Koenig, Eds. AAAI Press, 1467–1474.
- BREWKA, G., DELGRANDE, J., ROMERO, J. AND SCHAUB, T. 2015b. Implementing preferences with asprin. In *Proceedings of the 13th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'15)*, Calimeri, F., Ianni, G. and Truszczyński, M., Eds. Lecture Notes in Artificial Intelligence, vol. 9345. Springer-Verlag, 158–172.
- BREWKA, G., EITER, T. AND MCILRAITH, S., Eds. 2012. *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR'12)*. AAAI Press.
- CABALAR, P. AND SON, T., Eds. 2013. *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'13)*. Lecture Notes in Artificial Intelligence, vol. 8148. Springer-Verlag.
- CALIMERI, F., COZZA, S. AND IANNI, G. 2007. External sources of knowledge and value invention in logic programming. *Annals of Mathematics and Artificial Intelligence* 50, 3–4, 333–361.
- CALIMERI, F., FABER, W., GEBSE, M., IANNI, G., KAMINSKI, R., KRENNWALLNER, T., LEONE, N., RICCA, F. AND SCHAUB, T. 2012. ASP-Core-2: Input language format [online]. URL: <https://www.mat.unical.it/aspcomp2013/ASPstandardization>.
- CALIMERI, F., IANNI, G. AND TRUSZCZYŃSKI, M., Eds. 2015. *Proceedings of the 13th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'15)*. Lecture Notes in Artificial Intelligence, vol. 9345. Springer-Verlag.
- DE CAT, B., BOGAERTS, B., BRUYNOOGHE, M. AND DENECKER, M. 2014. Predicate logic as a modelling language: The IDP system. arXiv:1401.6312.
- DE MOURA, L. AND BJØRNER, N. 2008. Z3: An efficient SMT solver. In *Proc. of the 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'08)*, C. Ramakrishnan and J. Rehof, Eds. Lecture Notes in Computer Science, vol. 4963. Springer-Verlag, 337–340.
- DE POOTER, S., WITTOCX, J. AND DENECKER, M. 2013. A prototype of a knowledge-based programming environment. In *Proceedings of the 19th International Conference on Applications of Declarative Programming and Knowledge Management (INAP'11) and the 25th Workshop on Logic Programming (WLP'11)*, H. Tompits, S. Abreu, J. Oetsch, J. Pührer, D. Seipel, M. Umeda, and A. Wolf, Eds. Lecture Notes in Computer Science, vol. 7773. Springer-Verlag, 279–286.
- DELGRANDE, J. AND FABER, W., Eds. 2011. *Proceedings of the 11th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'11)*. Lecture Notes in Artificial Intelligence, vol. 6645. Springer-Verlag.
- DELGRANDE, J., SCHAUB, T., TOMPITS, H. AND WOLTRAN, S. 2008. Belief revision of logic programs under answer set semantics. In *Proc. of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, G. Brewka and J. Lang, Eds. AAAI Press, 411–421.
- DELGRANDE, J., SCHAUB, T., TOMPITS, H. AND WOLTRAN, S. 2009. Merging logic programs under answer set semantics. In *Proc. of the 25th International Conference on Logic Programming (ICLP'09)*, P. Hill and D. Warren, Eds. Lecture Notes in Computer Science, vol. 5649. Springer-Verlag, 160–174.
- DIMOPOULOS, Y., GEBSE, M., LÜHNE, P., ROMERO, J. AND SCHAUB, T. 2017. plasp 3: Towards effective ASP planning. In *Proceedings of the 14th International Conference on Logic*

- Programming and Nonmonotonic Reasoning (LPNMR'17)*, Balduccini, M. and Janhunen, T., Eds. Lecture Notes in Artificial Intelligence, vol. 10377. Springer-Verlag, 286–300.
- DODARO, C., GASTEIGER, P., LEONE, N., MUSITSCH, B., RICCA, F. AND SCHEKOTIHIN, K. 2016. Driving CDCL search. arXiv:1611.05190.
- DODARO, C., RICCA, F. AND SCHÜLLER, P. 2016. External propagators in wasp: Preliminary report. In *Proc. of the 23rd International Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion (RCRA'16)*, CEUR Workshop Proceedings, vol. 1745, 1–9.
- EÉN, N. AND SÖRENNSSON, N. 2003. Temporal induction by incremental SAT solving. *Electronic Notes in Theoretical Computer Science* 89, 4, 543–560.
- EÉN, N. AND SÖRENNSSON, N. 2004. An extensible SAT-solver. In *Proc. of the 6th International Conference on Theory and Applications of Satisfiability Testing (SAT'03)*, E. Giunchiglia and A. Tacchella, Eds. Lecture Notes in Computer Science, vol. 2919. Springer-Verlag, 502–518.
- EITER, T., FINK, M., KRENNWALLNER, T. AND REDL, C. 2012. Conflict-driven ASP solving with external sources. *Theory and Practice of Logic Programming* 12, 4–5, 659–679.
- FEBBRARO, O., LEONE, N., GRASSO, G. AND RICCA, F. 2012. JASP: A framework for integrating answer set programming with Java. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR'12)* BREWKA, G., Eiter, T. and McIlraith, S., Eds. AAAI Press, 541–551.
- FINK, M., GERMANO, S., IANNI, G., REDL, C. AND SCHÜLLER, P. 2013. ActHEX: Implementing HEX programs with action atoms. In *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'13)*, Cabalar, P. and Son, T., Eds. Lecture Notes in Artificial Intelligence, vol. 8148. Springer-Verlag, 317–322.
- FUSCÀ, D., GERMANO, S., ZANGARI, J., ANASTASIO, M., CALIMERI, F. AND PERRI, S. 2016. A framework for easing the development of applications embedding answer set programming. In *Proc. of the 18th International Symposium on Principles and Practice of Declarative Programming (PPDP'16)*, J. Cheney and G. Vidal, Eds. ACM Press, 38–49.
- GEBSER, M., GROTE, T., KAMINSKI, R., OBERMEIER, P., SABUNCU, O. AND SCHAUB, T. 2012. Stream reasoning with answer set programming: Preliminary report. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR'12)* BREWKA, G., Eiter, T. and McIlraith, S., Eds. AAAI Press, 613–617.
- GEBSER, M., GROTE, T., KAMINSKI, R. AND SCHAUB, T. 2011. Reactive answer set programming. In *Proceedings of the 11th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR '11)*. Delgrande, J. and Faber, W., Eds. Lecture Notes in Artificial Intelligence, vol. 6645. Springer-Verlag, 54–66.
- GEBSER, M., HARRISON, A., KAMINSKI, R., LIFSCHITZ, V. AND SCHAUB, T. 2015. Abstract Gringo. *Theory and Practice of Logic Programming* 15, 4–5, 449–463. URL: <http://arxiv.org/abs/1507.06576>.
- GEBSER, M., JOST, H., KAMINSKI, R., OBERMEIER, P., SABUNCU, O., SCHAUB, T. AND SCHNEIDER, M. 2013. Ricochet robots: A transverse ASP benchmark. In *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'13)*. Cabalar, P. and Son, T., Eds. Lecture Notes in Artificial Intelligence, vol. 8148. Springer-Verlag, 348–360.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B., LINDAUER, M., OSTROWSKI, M., ROMERO, J., SCHAUB, T. AND THIELE, S. 2015. *Potassco User Guide*, 2nd ed. University of Potsdam.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B., OSTROWSKI, M., SCHAUB, T. AND SCHNEIDER, M. 2011. Potassco: The Potsdam answer set solving collection. *AI Communications* 24, 2, 107–124.

- GEBSER, M., KAMINSKI, R., KAUFMANN, B., OSTROWSKI, M., SCHAUB, T. AND THIELE, S. 2008. Engineering an incremental ASP solver. In *Proc. of the 24th International Conference on Logic Programming (ICLP'08)*, M. Garcia de la Banda and E. Pontelli, Eds. Lecture Notes in Computer Science, vol. 5366. Springer-Verlag, 190–205.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B., OSTROWSKI, M., SCHAUB, T. AND WANKO, P. 2016. Theory solving made easy with clingo 5. In *Technical Communications of the 32nd International Conference on Logic Programming (ICLP'16)*, M. Carro and A. King, Eds., Open Access Series in Informatics (OASICs), vol. 52, 2:1–2:15.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2012. *Answer Set Solving in Practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool Publishers.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2014. *Clingo = ASP + control: Preliminary report*. In *Proc. of Technical Communications of the 30th International Conference on Logic Programming (ICLP'14)*, M. Leuschel and T. Schrijvers, Eds. Theory and Practice of Logic Programming, Online Supplement, vol. 14(4–5). URL: <http://arxiv.org/abs/1405.3694v1>.
- GEBSER, M., KAMINSKI, R., KÖNIG, A. AND SCHAUB, T. 2011. Advances in gringo series 3. In *Proceedings of the 11th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR '11)*. Delgrande, J. and Faber, W., Eds. Lecture Notes in Artificial Intelligence, vol. 6645. Springer-Verlag, 345–351.
- GEBSER, M., KAMINSKI, R., OBERMEIER, P. AND SCHAUB, T. 2015. Ricochet robots reloaded: A case-study in multi-shot ASP solving. In *Proc. of Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation: Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, T. Eiter, H. Strass, M. Truszczyński, and S. Woltran, Eds. Lecture Notes in Artificial Intelligence, vol. 9060. Springer-Verlag, 17–32.
- GEBSER, M., KAUFMANN, B., OTERO, R., ROMERO, J., SCHAUB, T. AND WANKO, P. 2013. Domain-specific heuristics in answer set programming. In *Proc. of the 27th National Conference on Artificial Intelligence (AAAI'13)*, M. desJardins and M. Littman, Eds. AAAI Press, 350–356.
- GEBSER, M., KAUFMANN, B. AND SCHAUB, T. 2013. Advanced conflict-driven disjunctive answer set solving. In *Proc. of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*, F. Rossi, Ed. IJCAI/AAAI Press, 912–918.
- GEBSER, M., OBERMEIER, P. AND SCHAUB, T. 2013. A system for interactive query answering with answer set programming. In *Proc. of the 6th Workshop on Answer Set Programming and Other Computing Paradigms (ASPOCP'13)*, M. Fink and Y. Lierler, Eds. Vol. abs/1312.6143. CoRR.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *Proc. of the 5th International Conference and Symposium of Logic Programming (ICLP'88)*, R. Kowalski and K. Bowen, Eds. MIT Press, 1070–1080.
- JANHUNEN, T., KAMINSKI, R., OSTROWSKI, M., SCHAUB, T., SCHELLHORN, S. AND WANKO, P. 2017. Clingo goes linear constraints over reals and integers. *Theory and Practice of Logic Programming* 17, 5–6, 872–888.
- KAMINSKI, R., SCHAUB, T. AND WANKO, P. 2017. A tutorial on hybrid answer set solving with clingo. In *Proc. of the 13th International Summer School of the Reasoning Web*, G. Ianni, D. Lembo, L. Bertossi, W. Faber, B. Glimm, G. Gottlob, and S. Staab, Eds. Lecture Notes in Computer Science, vol. 10370. Springer-Verlag, 167–203.
- KAUFMANN, B., LEONE, N., PERRI, S. AND SCHAUB, T. 2016. Grounding and solving in answer set programming. *AI Magazine* 37, 3, 25–32.
- LEONE, N., PFEIFER, G., FABER, W., EITER, T., GOTTLÖB, G., PERRI, S. AND SCARCELLO, F. 2006. The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic* 7, 3, 499–562.

- LIFSCHITZ, V. AND TURNER, H. 1994. Splitting a logic program. In *Proc. of the 11th International Conference on Logic Programming*. MIT Press, 23–37.
- NIEUWENHUIS, R., OLIVERAS, A. AND TINELLI, C. 2006. Solving SAT and SAT modulo theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(T). *Journal of the ACM* 53, 6, 937–977.
- OIKARINEN, E. AND JANHUNEN, T. 2006. Modular equivalence for normal logic programs. In *Proc. of the 17th European Conference on Artificial Intelligence (ECAI'06)*, G. Brewka, S. Coradeschi, A. Perini, and P. Traverso, Eds. IOS Press, 412–416.
- OSTROWSKI, M. AND SCHAUB, T. 2012. ASP modulo CSP: The clingcon system. *Theory and Practice of Logic Programming* 12, 4–5, 485–503.
- SABUNCU, O. AND LEITE, J. 2017. moviola: Interpreting dynamic logic programs via multi-shot answer set programming. In *Proceedings of the 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'17)*, Balduccini, M. and Janhunen, T., Eds. Lecture Notes in Artificial Intelligence, vol. 10377. Springer-Verlag, 336–342.
- SIMONS, P., NIEMELÄ, I. AND SOININEN, T. 2002. Extending and implementing the stable model semantics. *Artificial Intelligence* 138, 1–2, 181–234.
- SYRJÄNEN, T. 2001. Lparse 1.0 user's manual.
- ZHANG, Y. AND FOO, N. 2006. Solving logic program conflict through strong and weak forgettings. *Artificial Intelligence* 170, 8–9, 739–778.